The Wages of Nonemployment

Michael W. L. Elsby   Ryan Michaels   David Ratner

Matthew D. Shapiro[1]

February 15, 2018

Abstract

Male joblessness in the United States has risen significantly over the last half century, driven by an increase in the incidence of very long jobless spells and concentrated among the low-skilled. Motivated by this, we document the means by which those chronically out of work get by, how these means have evolved over time, and whether changes in the availability and generosity of nonwork benefits plausibly might have contributed to the rise of male joblessness. We find that the large rise in male nonemployment is challenging for canonical models of labor supply to explain given the empirical relationship between income and weeks work that we document. We explore the ability of refinements and extensions of canonical models to resolve this puzzle.

JEL codes: E24, J2, J3, J6. Keywords: X.

[1] Elsby: University of Edinburgh (mike.elsby@ed.ac.uk). Michaels: Federal Reserve Bank of Philadelphia (ryan.michaels@phil.frb.org). Ratner: Federal Reserve Board (david.d.ratner@frb.gov). Shapiro: University of Michigan (shapiro@umich.edu).

We are grateful to conference and seminar participants at St. Louis Federal Reserve and Census Bureau for helpful comments. We thank Katherine Richard and Paul van Vliet for excellent research assistance. All errors are our own. The views expressed here do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the staff and members of the Federal Reserve Board, or the Federal Reserve System as a whole.

Elsby is grateful for the hospitality of the Faculty of Economics and INET Institute at the University of Cambridge where part of this work was completed. Elsby, Michaels and Shapiro gratefully acknowledge financial support from the UK Economic and Social Research Council (ESRC), Award reference ES/L009633/1.

Over the last half century, rates of joblessness among prime-aged U.S. men have exhibited a significant upward trend. Nonemployment rates for this group have doubled, driven by rises in the incidence of very long jobless spells. The increase in extended jobless spells is particularly pronounced among the non-college educated, 15 percent of whom now do not work at all in a year. What accounts for this large, persistent decline in labor market attachment among a subgroup that was once a mainstay of the U.S. labor market? By what means have these men subsisted in an economy not traditionally known for the availability of long-term non-work benefits? And how have non-work income and nonemployment rates among U.S. men evolved together? These are the questions this paper seeks to address.

After revisiting and updating a set of stylized facts on male joblessness in section 1, we explore in section 2 empirical evidence on the distribution of income by weeks worked, its level, composition, and evolution over time. Strikingly, microdata from the March Current Population Survey paint a picture of *stasis* in the distribution of income by weeks worked, a fact that can be difficult for workhorse labor supply theory to square with sharply declining employment. While the household income of workless men is nontrivial—providing a perspective on how they get by in the absence of earned income—both this level, *and its relation to the income of working men*, are similar over time. From this point of view, the decline in employment has not been accompanied by any clear increase in the relative (pecuniary) payoff to nonworking time. By the same token, although the average composition of income is intuitive—for example, unemployment insurance replaces income for part-year jobless spells, and Social Security contributes to income among the chronically jobless— there is, again, little evidence in March CPS data to suggest that this has changed radically over time.

CPS microdata are known to have disadvantages in measuring certain types of income. For example, the CPS does not report tax credits such as the EITC. In addition, it is well known that there is significant underreporting of transfer income in the CPS (Meyer et al, 2016). We expand on CPS data through a series of extensions in order to capture important missing pieces of the CPS income landscape; namely we account for net tax liability using NBER's TAXSIM programs; for public and private health insurance; and underreporting of AFDC, Food Stamps, SSI, and other programs using the TRIM model, a long-standing microsimulation program maintained by the Urban Institute.

Considering a more expansive measure of income shows that, in contrast to the raw CPS data, incomes of those who worked less of the year was buoyed somewhat relative to incomes of the fully employed. In other words, the estimated *gradient* of income with respect to weeks worked—a statistic that can be indicative, as we show, of the return on labor supply—did decline modestly on net over the last 50 years. Looking a little more closely, though, reveals that the comovement of the income gradient with nonemployment is still surprisingly weak. Decades in which employment declines most precipitously are often unaccompanied by meaningful declines in the gradient, whereas epochs in which the gradient declines lack substantial changes in employment.

One challenge in interpreting these findings, of course, is that the gradient can only measure the payoff from working using the wage incomes of those who have *chosen* to work. However, if public assistance increases or real wages fall, the men who continue working are likely the ones who face especially high returns from doing so. As a result, comparing their income against the incomes of the (newly) nonworking may not imply much of a decline in gradient. There are few avenues one can take with the data to mitigate this problem.[2]

A complementary exercise approaches the issue from the perspective of theory, and asks if this selection mechanism can, within a canonical theoretical model, account for both the fall in employment and relative stasis in the gradient. In section 3, we conduct such an exercise. In canonical models with balanced-growth preferences, proportionate changes in wages and nonwork income have exactly offsetting effects on labor supply. Accordingly, individual labor supply changes in the model reflect changes in the *ratio* of non-work income to in-work wages, or the *replacement rate*. The higher is this replacement rate, the lower will be his labor supply. It follows that broad-based rises in aggregate nonemployment have their origins in widespread rises in replacement rates, under standard labor supply theory.

A key question is how a change in the average replacement rate (a structural parameter) shapes the observed reduced-form relationship between income and weeks worked that is estimable with suitable microdata. This mapping is especially intuitive in the absence of selection, that is, in the special case where each individual faces the same replacement rate. In that case, individuals face better non-work benefits to in-work wages, which discourages labor supply, and weakens its association with income. As a result, the weeks-income profile becomes more shallow. In a model with heterogeneous replacement rates, there is a potentially important selection mechanism in which only relatively low-replacement-rate workers continue to work in the face of an increase in the average rate. We assess this mechanism quantitatively within a realistically calibrated version of the model. Even here, we find that an increase in the average replacement rate induces a decline in the gradient that is counterfactually large relative to the change in employment outcomes.

1. A picture of male joblessness, 1967 to 2015

[2] For instance, one can estimate the gradient only using data on men who work less than, say, 26 weeks, with the idea being that these men are more likely to face similar payoffs from working. One can also attempt to control for unobserved effects by linking respondents' outcomes in the March CPS across years and estimating the gradient including individual fixed effects. We have carried out both exercises; the bottom line of our empirical findings is unaffected.

Male joblessness in the United States has risen significantly over the last half century. Here, we revisit and update the empirical evidence for this phenomenon, which was first noted in the early work of Juhn, Murphy and Topel (1991, 2002), and studied more recently in Elsby and Shapiro (2012). Recently, Aguiar et al. (2017) examines how technological advances may have enhanced the value of leisure and reduced labor supply of younger cohorts. Krueger (2016) explores the role of rising work-limiting disabilities and health problems among low-skilled men in their declining employment.

Figure 1 summarizes the evidence. It uses data from the March Current Population Surveys (CPS) from 1968 to 2016. Each March CPS asks respondents to report their weeks worked over the prior calendar year.[3] Following Juhn et al., we use these data to measure the fraction of the year spent out of work, defined as weeks not worked divided by 52. Aggregating across respondents, and weighting by the CPS sampling weights, provides us with a time series for the nonemployment rate for the calendar years 1967 through 2015.

We restrict our samples to men, aged 25 to 54, who report that they are not in school, retired, in the military, or self-employed (in the longest job held in the prior year). We additionally focus on a subset of men that we refer to as "primary males" in a household; this set includes a male household head, and a spouse or cohabiting male partner of a female household head.[4] We focus on this subsample because, as we shall see, such prime-aged men historically have displayed a high degree of attachment to the labor market.[5]

Figures 1A and 1B plot the aggregate trends. Prime-aged male nonemployment in our sample has doubled as a trend phenomenon over the last fifty years, varying from around 6 percent in the early 1970s, to about 12.5 percent since the turn of the century. This trend is driven by a rise in the incidence of very long jobless spells. Full-year nonemployment rates have risen fourfold over the same period, from approximately 2 percent to around 8 percent (Figure 1A). And, as one would expect, much of the trend rise in nonemployment can be traced to rises in nonparticipation (rather than unemployment), which has more than doubled from around 4 percent in the early 1970s to around 9 percent in recent years (Figure 1B).

---

[3] Between 1968 and 1975, the March CPS records a respondent's weeks worked only in discrete categories, binned as follows: 1-13, 14-26, 27-39, 40-47, 48-49, or 50-52 weeks. Following Juhn et al., we assign the weighted average weeks worked in each category from the 1976-1990 March CPS samples to the 1968-1990 observations.

[4] We identify cohabiting partners in a similar way to Casper and Cohen's (2000) "Adjusted POSSLQ" method. We include households in which an adult man is living with an unrelated, unmarried female household head, and in which there are no additional unrelated adults, with the exception of children of the cohabiting male. This is intended to eliminate households in which multiple adult male roommates are living with one adult women, and to retain cases in which a cohabiting couple is living together with the man's children, who are unrelated to the female household head.

[5] These sample restrictions differ somewhat to those used by Juhn et al., who focus on white men with 1-30 years of potential labor market experience. The trends are not materially affected by these differences.

Figure 1C reveals that these aggregate trends have been driven in large part by a steep skill gradient to the rise in male nonemployment, reiterating a further theme of earlier literature. Among prime-aged men with a high school education or less, nonemployment rates have tripled since the early 1970s, from around 7 percent to around 20 percent in Figure 1C. Although their college-educated counterparts also are working less than in the past, the incidence of joblessness is much lower among the more-skilled. For this reason, the focus of the remainder of the paper will be to investigate more deeply the outcomes of noncollege-educated men in our samples.

Although not the focus of our ensuing analysis, in the remaining panel of Figure 1 we highlight the enduringly alarming racial gradient to the rise in male joblessness. Because white men account for the vast majority of the male population, their nonemployment rates by education resemble those in Figure 1C. But, the trends for black men paint an even more startling picture. Figure 1D reveals that more than a quarter of noncollege-educated prime-aged black men have been out of work over the last two decades, around double the nonemployment rates of similarly-educated white men.

The message of Figure 1, then, is that rates of male joblessness have trended upward significantly in the United States over the last half century, that these trends have occurred even among prime-aged men who traditionally were highly attached to the labor market, have been accompanied by rises in persistent, often full-year spells, and withdrawal from the labor force, and have been borne especially by low-skilled and black men.

A host of questions emerge naturally in the light of these facts. By what means do those chronically out of work get by, particularly in an economy not known for the widespread availability of long-term nonwork benefits? And how have these means evolved over time as the incidence of nonemployment has risen? Have changes in the availability and generosity of nonwork benefits contributed to the evolution of male joblessness? These questions are the focus of each of the ensuing sections.

2. The wages of nonemployment

Motivated by these questions, in this section we document the level, distribution, composition by source, and evolution over time of income by weeks worked in the United States. We present two sets of empirical results, described in each of the following subsections. In the first, we return to the March CPS data samples for the calendar years 1967 through to 2015 used in section 1 to explore the detailed survey measures of household income that they provide. In the second, we augment these with additional estimates of public taxes and benefits that extend and refine the March CPS measures. Throughout, we focus on men with a high school diploma (or equivalent) or less, reflecting the concentration of the rise in nonemployment among this group documented in Figure 1C. All other sample restrictions remain as described in section 1.

2.1 A view from the March Current Population Survey

In our first set of results, we use the individual-level detail in the March CPS microdata to create a consistent decomposition of household income over the sample period. We begin by dividing income into that received by the Primary Male (as defined above), that received by a female spouse or cohabiting partner (if one is present), and that received by other household members (any children, non-spouses, non-child relatives, or non-relatives).[6]

Within the household head's income, we further distinguish between earned and unearned income. Since we restrict our samples to men who are not self-employed, household heads' earned income is comprised entirely by wage and salary income. Unearned income, by contrast, has a richer detail. We distinguish between income from public assistance (welfare), unemployment insurance (and related benefits), Social Security (disability insurance), and other income (which includes interest, dividends, rent, alimony, child support, contributions from friends, and so on).[7] Table 1 provides a summary and further detail on the decomposition of household income.[8] All income variables are deflated to $2015 using the CPI-U price index.

Figure 2 provides an illustration of the distribution of household income by weeks worked of the household head over time. Weeks worked are binned into quarters of the year, 0-12 weeks, 13-25 weeks, 26-38 weeks, and 39-52 weeks. To summarize long-run trend movements, we plot the distribution of income conditional on weeks worked for two sub-periods, 1967 to 1989, and 1990 to 2015, respectively. Recall from Figure 1C that nonemployment rates among noncollege-educated prime-aged men averaged 12 percent in the earlier sub-period, and 17 percent in the later sub-period.

Levels. A first message of Figure 2 is that, even in households with workless primary males, the average levels of household income faced by low-skilled prime-aged men are nontrivial. For households in which the Primary Male worked 12 weeks or fewer in the preceding year, annual household income averaged approximately $30,000. This stands at a little over 40 percent of the average income faced by households in which the male worked in excess of 39 weeks in the prior year.

Second, a striking feature of Figure 2 is the near-stasis over time of mean household income by the male's weeks worked. Differences in mean outcomes pre- versus post-1990 are hard to discern. For example, the ratio of household income in the last bin of weeks

[6] It is also possible to separate income sources by relationship to household head—for example, children's wage and salaries, and so on. We do not delve into that degree of detail.

[7] Inconsistencies in how some variables are coded over time requires aggregation of some income categories. We have tried to devise the most detailed but consistent breakdown.

[8] Building household income from the bottom up can lead to cases in which total household income (the sum across individual incomes for each income type) does not sum to the household income variable created by the CPS. This can happen, for example, if one of the sub-income categories is top-coded for a person, but total household income was not top-coded by the CPS.

worked to that in the first bin rises only from 42 percent to 44 percent. Thus, throughout the 1967-2015 sample period, households with chronically workless men appear to not be falling much further behind more attached man, at least when viewed on the average.

Composition. Figures 2A and 2B illustrate the average composition of household income by source, and thereby provide a perspective on the means by which workless men get by. As men's weeks worked decline, and wage and salary income recedes, three forms of replacement income appear to play a role, and in an intuitive way.

First, unemployment income is increasingly received by men with part-year nonemployment spells, but maxes out at approximately $3,500 to $4,000 per year for those with very few weeks worked.[9] This accords with the typical durations of unemployment insurance claims of 26 weeks (outside of recessions).

Second, in addition to this, workless men with the fewest weeks worked receive, on average, a little over $3,500 per year of replacement income from Social Security benefits. Consistent with the nature of these benefits, this income source is trivial for men with part-year nonemployment spells.

Third, men increasingly receive replacement income in the form of Public Assistance (welfare) benefits rising to around $2,500 per year for men who worked 12 weeks or fewer in the prior year.

In aggregate, these forms of replacement income contribute considerably to the income of households with workless men, summing to approximately $10,000 per year on average for men in the lowest weeks worked category.

Figure 2 highlights a further prominent source of household income coming from spouses, cohabiting partners and other household members. This has three features. First, it contributes significantly to overall household income: Income from spouses or cohabiting partners consistently contributes in excess of $10,000 per year on average for all categories of the male's weeks worked; income from other household members contributes another $5-7,000.

Second, the average level of income from spouses, partners and other household members is strikingly invariant with respect to the primary male's weeks worked. The one notable exception is that, in the later post-1990 period, such income is around $4,000 larger per year in households whose primary males worked 39 weeks or more.

Third, income from spouses, partners and other household members has risen over time. Spousal or partner income averaged around $3-6,000 more per year in the later period; other household member income rose by around $1,500.

Interestingly, the latter more than offset the declines in income faced by the primary male in the household, which have been driven especially by declines in their wage and

[9] This figure is not conditional upon receiving that type of transfer, but rather average across all households in the sample.

salary income, and account for the modest rise over time in overall household income across all weeks worked categories noted earlier.

Distribution. Up to this point, we have summarized *mean* levels of household income by primary male weeks worked, which have painted a picture of near-invariance over time. But there have been important trends in the conditional *distribution* of income. Figures 2C and 2D plot the mean, median and interquartile range of household income for each bin of the male's weeks worked. These reveal three findings.

First, although average incomes in households with workless men are nontrivial, there is considerable heterogeneity in income across these households. Most notably, the lower tail among households whose primary male worked fewer than 12 weeks encompasses remarkably low incomes, with the lower quartile receiving as little as $12-13,000 per year.

Second, this picture of considerable inequality in income is mirrored at all levels of the primary male's weeks worked. Prior to 1990, the interquartile range of household income was equal to 107 log points in the lowest bin of weeks worked, and successively 92, 79 and 61 log points in each consecutively higher weeks-worked bin.

Third, while there has been near-stasis in mean and median incomes by weeks worked, there has been a large rise in income inequality. But, again, this has been remarkably uniform across the distribution of men's weeks worked: Rises in income inequality in households in which the primary male works most of the year have been mirrored by rises in inequality in households with mostly workless men. Interquartile ranges across all weeks-worked bins rose by approximately 11-19 log points post-1990 versus pre-1990. Interestingly, the well-documented rise in wage inequality among those in work has been mirrored by a near-symmetric increase in inequality in nonwork or other sources of income among those out of work.

These distributional outcomes underscore that, while average levels of income among the growing set of households with jobless men have remained nontrivial over time, there are large and growing differences in incomes across them, with an increasing fraction of households reporting very low annual incomes.

Evolution over time of the income gradient. A key motivation for our present analysis of the household incomes by the weeks worked of noncollege-educated men is the large rise over time in the nonemployment rates faced by this group. An important question, then, is whether the distribution of household income by primary male weeks worked has also varied over time; and, in particular, whether it has evolved in a way that mirrors the trends in male joblessness. In other words, as employment has declined, do the data paint a picture of declining returns to work viewed through the lens of the relationship between weeks worked and incomes?

Figure 2 provides a coarse sense of this, suggesting little change before versus after 1990. We now enrich that analysis in two ways. First, we explore in more detail the time profile of income by weeks worked over the whole sample period, not just over the two broad subperiods. Second, the pattern of income by male weeks worked in Figure 2 is likely to be distorted by compositional differences between households in which the man works different numbers of weeks. For this reason, we investigate the extent to which differences in attributes between these households alter the level and time profile of income by primary male weeks worked.

To do so, we estimate a simple summary statistic for the relationship between household income and the fraction of the year worked by the primary male. Specifically, for each year of available data t we estimate least-squares regressions of the form

$$\ln y_{it} = a_t + g_t h_{it} + \psi_t x_{it} + e_{it}, \tag{1}$$

where y_{it} is household income, h_{it} the fraction of the year worked by the primary male (weeks worked divided by 52), and x_{it} a vector of controls that we will discuss shortly.

The key coefficient of interest is g_t, which provides an estimate of the semi-elasticity of income with respect to the fraction of the year worked by the primary male, what we shall term the *income gradient*. g_t is thus a simple summary statistic for the relationship between a household's income and the weeks worked of its primary male.[10]

Figure 3 plots time series for these estimates of the income gradient for our sample of men with a high school degree or less. It plots two series. First, as a benchmark, the "no controls" series simply plots the income gradient estimated purely from cross-sectional differences in household income by head's weeks worked. In this way, it is a year-by-year analogue to the picture presented in Figure 2.

The evolution of the income gradient in Figure 3 provides an interesting contrast with the path of male joblessness in Figure 1. Four broad eras can be identified. First, over the course of the 1970s and early 1980s, the noncollege male nonemployment rate doubled from around 7 percent to 15 percent. Strikingly, this was accompanied by a significant *rise* in the income gradient. At the start of the sample, a household in which the primary male was full-year nonemployed ($h = 0$) faced household income on average 110 log points lower than a household in which the primary male worked full-year ($h = 1$). By the mid-1980s, this difference had grown to around 120 log points: The larger pool of households with workless males faced relative incomes that were even lower by the mid-1980s.

A second era runs from the mid-1980s through to the early 1990s. Over this period, the income gradient reverses course, falling quite sharply to restore a near-110-log-point difference. Yet, this period witnessed near-stasis in the male nonemployment rate. That is, as the relative household incomes of households with workless primary males recovered, over just a handful of years, the incidence of male joblessness barely moved.

[10] We have explored different functional forms for the relationship between log household income and male weeks worked, and have found that the linear specification in (1) provides a reasonable fit. The results are robust to specifying equation (1) in the level of income, or more non-parametric approaches including Lowess estimates, or using variation in weeks worked only in the poles of the weeks worked distribution. Results are available upon request.

A third era, spanning the early-to-mid 1990s through to the period preceding the Great Recession, exhibits relatively modest changes in both male nonemployment, which remained around 15 percent, and the income gradient, which remained close to 110 log points.

Finally, the Great Recession was accompanied by large rises in noncollege male nonemployment, up to 23 percent at its peak in 2010, and by a modest fall in the income gradient, down to around 105 log points over the period. Of the four historical episodes, this provides one example over which the pool of households with workless males has risen at the same time as their relative incomes have grown.

That episode aside, the broad picture painted by Figure 3 has two themes. First, at medium-run frequencies, both the male nonemployment rate and the income gradient exhibit nontrivial variation. But the medium-run comovement between these series is hard even to sign. Second, viewed over the long run, there is little relationship between male joblessness and the relative incomes their households face. The approximate tripling of the incidence of noncollege primary male nonemployment over the span of the sample is associated with essentially no change in the income gradient.

This impression is strengthened as one adds demographic controls. The second series in Figure 3 plots income gradients based on estimating (1) with controls x_{it} that include a quartic in age, a quartic in potential labor market experience, marriage and race dummies, the number of household members, and the number of household members less than 18 years of age.[11] Adding these controls reduces estimated income gradients by approximately 5-10 log points, indicating that some of the differences in income by weeks worked are accounted for by selection on these observables. But the evolution over time of these estimates remains very similar to that depicted by the series without controls.[12]

2.2 Extensions and refinements of the March CPS measures

The March CPS has the virtue of including substantial detail on household income over a long sample period that covers the rise in male joblessness. In this subsection, we seek

[11] These controls yield adjustments for household composition that are similar to those implied by related "equivalence scales" applied to official poverty thresholds by the Census Bureau. For instance, our estimates suggest that, relative to a single individual, a married couple has 27 percent higher income. This compares to a 29 percent increase in the poverty threshold for a two-person household relative to a one-person household. See http://www.census.gov/data/tables/time-series/demo/income-poverty/historical-poverty-thresholds.html. Further, our estimates are robust to adjusting household income by the square root of the number of household members as recommended by poverty researchers at the OECD.

[12] We also use the rotation structure of CPS interviews to link respondents longitudinally across successive March CPS. This allows us to construct a sequence of two-year individual panels for the majority of years in the sample. We use the two-year samples to estimate income gradients according to a first-differences version of equation (1); rather than using cross-sectional variation in weeks worked and incomes, these estimates rely on within-individual changes in those variables. The upshot is much the same. While the income gradients are lower on average, indicating selection of individuals across weeks worked in the purely cross-sectional results, the income gradient actual rises substantially over time.

to address potential drawbacks of these data, and their likely implications for the evolution of household income by primary male weeks worked. In particular, we address three areas of concern with respect to the raw CPS measures: the omission of the (progressive) income tax system, its accounting for health insurance, and the under-reporting of benefit income. We describe in what follows a series of imputation procedures used to adjust for these shortcomings; further details on these procedures are provided in the Appendix.

After-tax income. Up to this point, we have been using CPS measures of household income, which ask respondents about their income gross of tax. Of course, what matters to the household, both from the point of view of welfare, and for the decision to work, is the profile of *after-tax* income.

Accordingly, we have extended our preceding analysis by estimating after-tax income of households in the CPS. For each household in the CPS, we estimate a net federal tax liability (taxes paid minus credits received, including the EITC) using the National Bureau of Economic Research's (NBER) simulation program, TAXSIM. Drawing on a repository of federal tax law, and using CPS data on income sources, family and household composition, and other variables, TAXSIM provides an estimate of each household's income tax liability, defined as the sum of federal income taxes and the employee-share of FICA payments.

Accounting for federal taxes does not appear to alter dramatically our preceding findings. As shown in Figure 4, the profile of income by weeks worked is, on average, shallower: Progressivity in the tax and benefit system reduces the return to high weeks of work. In addition, taking account of taxes does imply less of an increase in the gradient during the first 15 years of the sample from 1967 to 1981, consistent with relatively larger increases in marginal tax rates for higher-income individuals than for lower (Tax Policy Center, 2015). Changes in taxation during the 1980s and 1990s, such as the introduction of federal taxes on UI and Social Security income and the expansion of the Earned Income Tax Credit, are captured by TAXSIM but appear to leave less of an imprint on the estimated income gradients. Overall, the trends in the income gradient are little changed by taxes.[13]

Accounting for health insurance. An additional shortcoming of our baseline CPS estimates is that they do not account for health insurance as an implicit source of income. Yet changes in the provision of public and private health insurance could have important implications for the evolution over time of effective household income by weeks worked.

[13] The tax liability of SSDI recipients in particular was limited by the progressive structure of the tax. The expansion of EITC also improved incomes not just for those working the fewest weeks but also those working part-year and particularly low-wage workers that worked any fraction of the year; thus, the EITC should not be expected to decrease the gradient substantially as we find.

Access to publicly funded health insurance has increased in recent decades. A 1973 law extended Medicare to SSDI recipients after two years in the disability insurance program. Since 1987, a series of laws has broadened eligibility for Medicaid—once largely restricted to single-parent households on welfare (AFDC)—to all households below an income threshold.[14] Notably, the expansion of social insurance coincided with a decline in the incidence of employer coverage.[15]

Since 1980, the March CPS has asked respondents if they have health insurance, but not about the value of their plans. Rather, for respondents covered under employer-provided plans, the Census provides imputed values for the employer's contribution in the March CPS. In the same vein, it has typically provided imputations for the "market value" of Medicaid and Medicare. Census' estimates are unavailable in the CPS prior to 1992, and after 2011 for Medicaid and Medicare.[16] We apply a simple imputation strategy to span a longer period than afforded by the estimates included in the CPS microdata.

For the value of employer-provided health insurance, we use data on benefit costs from the Bureau of Labor Statistics' Employer Costs of Employee Compensation (ECEC) program, which is derived from the National Compensation Survey. For a broad set of occupations within goods- and service-providing industries, the ECEC data enable us to calculate the ratio of the average employer cost of health insurance (i.e., the employer contributions to plan premium) to average wages. We apply this ratio to scale up earnings of each CPS respondent in the corresponding industry-occupation cell.[17] These ECEC data are available from 1988 forward.

Among nondisabled respondents who report Medicaid receipt, we impute a value of the program based on Medicaid expenditure per beneficiary. State-level expenditure data are available by broad age group and, separately, for the disabled, whom we identify in the CPS as SSI recipients; SSI receipt is an automatic qualifier for Medicaid. Our data on Medicaid extend from 1987 to 2014, which encompasses the period during which Medicaid was extended to non-AFDC families. We assume that individuals value Medicaid at 50

[14] See Gruber (2000) for a discussion of the legislative history up through the 1990s. The Affordable Care Act (ACA) further lifted the threshold applied to children aged 6 to 18, and enabled states to extend Medicaid to adults with no dependents whose income is less than 133 percent of the federal poverty line. However, the Medicaid expansion only went into effect in 2014, very late in our sample. In ongoing work, we are attempted to correct for changes in Medicaid after the ACA.

[15] The March CPS indicates that, among households whose prime-age male primary male had no college experience, the share covered by an employer plan fell from 75 percent in 1987-89 to 64 percent in 2014-16.

[16] Census' imputations of employers' contributions and the market value of Medicaid are first available in 1992, but the value of Medicare is not available until 2004.

[17] The Census' imputation strategy uses microdata from the 1977 National Medical Care Expenditure Survey (NMCES) to estimate the relationship between certain observables and employer contributions. Though this imputation can be carried out beginning in 1980, when the CPS first records employer coverage, the Census only provides the imputations since 1992. In ongoing work, and as a complement to our present strategy, we have replicated the Census regression specification on NMCES data and will impute contributions from 1980 to now.

cents on the dollar, following estimates in Finkelstein, Mahoney, and Notowidigdo (2018).[18]

Further, we assign a value of Medicare to each self-reported non-elderly Medicare recipient equal to average Medicare expenditure per disabled enrollee. Since the CPS first asked about Medicare receipt in 1980, we impute based on Medicare data for the years 1979 to 2013.[19]

Incorporating health insurance has a notable impact on the estimated income gradients. For example, incorporating the value of Medicare reduces the income gradient by about 5 log points on impact in calendar year 1979, the first year of the imputation. Likewise, incorporating Medicaid benefits also reduces the gradient by about 6 log points in 1987, as non-working individuals are more likely to receive these benefits than others. However, private insurance raises the gradient by almost the same degree in 1987 because, of course, those receiving private insurance are very likely to be working. Additionally, the estimated gradient declines more sharply in the latter half of the 1980s and the first half of the 1990s. This appears to reflect an increase in take-up of SSDI, and the associated Medicare benefit, as well as the phase-in of Medicaid coverage of older (age 6-18) children.

These estimates provide a first sense of the potential role of health insurance in household resources. However, they nonetheless are likely to underestimate the role of health insurance. Recall that our imputations are allocated only to respondents who report insurance in the CPS microdata. Yet there is evidence that CPS respondents under-report both private and public insurance coverage (see, for example, Peterson and Devere 2002). A more complete treatment of this under-reporting is beyond the scope of our analysis. However, in the case of certain public benefits, we can make relatively straightforward adjustments for under-reporting, a point to which we now turn.

Under-reporting of benefit income. The receipt of certain public benefits is underreported in several household surveys, including the CPS. Moreover, the extent of this under-reporting appears to have increased over time. For instance, comparing CPS estimates against administrative data sources, Meyer, Mok, and Sullivan (2015) find that the share of SNAP/Food Stamp expenditures captured by the CPS has declined by 0.6

[18] There is a voluminous literature on the question of how to monetarily value in-kind transfer programs. Many of these issues are covered in the exchange among Ellwood and Summers (1985), Blinder (1985), and Rees (1985).

[19] Medicaid and Medicare expenditure per enrollee is a practical analogue to employers' contributions. Consider the simple, but instructive case where employers pay the full premium, and private insurance is actuarially fair. The employer's contribution then equals the expected value of a claim. Meanwhile, total expenditure on public insurance is the product of the share s of enrollees who file claims; the total number of enrollees; and expenditure per claim, x. Thus, expenditure per enrollee is $s \cdot x$, which equals the expected value of a claim. Of course, this analogy neglects that compensation is fungible—employees can, in principle, choose between benefits and wages—to an extent that public assistance is not. As a result, public expenditure is likely to overstate recipients' valuation of the benefit, which is why the former is discounted based on Finkelstein et al estimates.

percentage points per year since 1980, and that the CPS currently understates these benefits by as much as 42 percent.

We carry out two exercises to explore the sensitivity of our results to the rise in under-reporting. First, certain benefits can be imputed to CPS respondents by assessing the probability of participation in light of the statutory parameters of the programs. Second, we can re-run our analysis using data from the Survey of Income and Program Participation (SIPP), which was developed to measure eligibility for and participation in public assistance programs. SIPP is available only since 1984, though. For the sake of brevity, SIPP results are reported in the Appendix.

We impute income from certain benefit programs using the Urban Institute's longstanding simulation program known as the Transfer Income Model or TRIM.[20] Specifically, we apply TRIM to impute income from the following four programs: AFDC/TANF (welfare); SNAP/Food Stamps; Supplemental Security Income (SSI); and housing assistance including, for example, Section 8 vouchers. TRIM's estimates are consistently available for each program starting in 1994.[21]

The imputation involves two steps: First, it determines a CPS respondent's eligibility; and second, it assigns a benefit payment based on the probability of take-up. For all four programs, the eligibility criteria are clearly articulated in federal and state statutes. TRIM then draws on auxiliary data where necessary to determine if a household qualifies for one (or more) of the benefits. For instance, one of the income tests for SNAP cannot, in general, be evaluated based on March CPS data alone—for example, income net of certain expenditures, such as child care and out-of-pocket medical expenses, must not exceed the federal poverty line. TRIM estimates these expenditures for CPS households based on Department of Agriculture surveys of food stamp participants (Giannarelli 1992). Likewise, for each benefit program, a probability of take-up is estimated based on supplementary data and assigned to each eligible CPS household that does not report receipt (eligible households that do report receipt are assigned a probability of one).[22] The probabilities are then scaled so that the implied number of participants in each state equals the count derived from administrative data.

Among the benefits that are not consistently included in TRIM's simulation program is unemployment insurance (UI). We devise a simple imputation method for UI that is described in full in the Appendix. In brief, we use CPS microdata to estimate the

[20] The current vintage of the microsimulation model is known as TRIM III. For several decades, the Department of Health and Human Services and related agencies (e.g., the former Office for Economic Opportunity) have financially supported the TRIM program and used it for policy analysis. See Zedlewski and Giannarelli (2015) for an overview and history of the TRIM program.

[21] TRIM publishes occasional studies of Medicaid take-up; for instance, on take-up among children, see Finegold and Giannarelli (2014). However, TRIM does not consistently provide imputed estimates of Medicaid for the full sample of CPS respondents.

[22] In the case of SNAP, the probability of take-up is estimated in TRIM based on observed participation in the SIPP. On the challenges of integrating the SIPP more fully into TRIM, see Giannarelli (1992).

relationship between average weekly UI receipt (conditional on receipt) and average weekly earnings (conditional on working) in each state. We then use this mapping to impute UI income to each CPS respondent who reports any weeks of layoff during the calendar year.[23]

As expected, incorporating the TRIM estimates shifts the estimated gradient down at all points in time, since the TRIM adjustments pick up a higher share of (true) benefit income. Moreover, the estimates including TRIM-simulated benefits from 1994 to 2014 decline by about 5 log points more than in the raw CPS data without the TRIM adjustments, reflecting the increase in benefit under-reporting in the CPS.

Putting it all together. In Figure 5 we combine all of the extensions and imputations of the CPS into a consistent series for the income gradient (from 1967-2013 through which we had a Medicaid adjustment). The red line repeats the series from Figure 4 using the after-tax gradient and controlling for demographic variables. We then combine the series implied by imputations for benefit programs (Food Stamps, AFDC, SSI, and housing subsidies) as well as our imputations for public and private health insurance benefits.[24] Importantly, the new series does indeed decline from the beginning of our sample to the end, from about 0.75 in 1967 to 0.65 in 2013. Despite the more pronounced downward tilt in the gradient, it is not apparent that the time series movements line up well with the increase in nonemployment with any regularity. For example, from 1967 to about 1985, the gradient increased by about 10 log points while the full-year nonemployment rate *increased* by about 5 percentage points. Moving forward, while the full-year nonemplyment rate continued its upward march over time, the gradient actually fed substantially from the mid-1980s to the early 1990s. We will return to the estimated income gradient below in a quantitative assessment of a standard labor supply model.

2.3 Revisiting trends in household earnings and benefits

Recall that changes in the income gradient reflect changes in the primary male's wages, public assistance, and other household members' earnings. Each of these

[23] This is of course an upper bound on UI income, since many of the unemployed are either ineligible or do not take up. As a last step, we want to sample from this potential universe of UI recipients so that total imputed UI income in the CPS matches that reported in administrative data. The latter work is underway.

[24] We do this by "backcasting" the imputed series for years they aren't available in the following manner. Take Medicaid as an example, which reduces the gradient by about 6 percentage points on impact in 1987. We use the ratio of real per capita Medicaid costs as a share of those in 1987 to backcast the contribution to the gradient in each year from 1967 to 1986. For example, in 1972, Medicaid costs were about half of those in 1987, so the contribution to the gradient is about -3 percentage points in 1972, half of the effect in 1987. We do a similar procedure for Medicare and Food Stamps. We use the employer costs for employee health insurance plans from NIPA data to backcast the contribution of private insurance on the gradient.

components has been studied extensively in the literature. In this section, we review each of them in the light of the preceding evidence on the evolution of the income gradient.

One of the most extensively documented facts of the U.S. labor market is the decline in real wages among noncollege educated men.[25] Over the last five decades, March CPS data reveal that real weekly earnings of primary males with no college experience have fallen almost 15 percent. It would seem to follow that these men face a substantially lower income gradient: when the wage is lower, a week less of work implies a smaller decline in income.

However, several public assistance programs tie benefit payments to wages to one degree or another. Hence, benefits will typically decline, too, as wages fall, which at least partially offsets the effect of lower real wages on the income gradient.

The clearest example of this is Unemployment Insurance, or UI. A claimant's weekly UI payment replaces a fraction of his average weekly earnings measured at some point within the year preceding his unemployment spell. The Employment and Training Administration supplies estimates of UI replacement rates for the years 1997-2017 using data on prior earnings of UI claimants. The replacement rate began and ended this period at about 46 percent.[26] For earlier years, replacement rates can be estimated using reports of earnings of all UI-eligible workers in the ETA's Handbook 394 data. Again, these data show no sustained rise in the UI replacement rate.[27]

Other public assistance programs also anchor benefits to an individual's own prior wages, albeit to a lesser degree. In the case of SSDI, for instance, the annual benefit is an increasing concave function of the recipient's "indexed" average annual lifetime earnings, which is derived by inflating one's own past earnings according to growth in *average* earnings (Autor and Duggan 2003 and 2006). This inflation factor has become more substantial as earnings inequality has widened. For a middle-aged male at the lower end of the earnings distribution, the share of annual earnings that can be replaced by SSDI has risen by 7 percentage points since 1984 (Muller 2008).[28] In the case of food stamps,

[25] For an updated analysis of the trends in real wages, see Acemoglu and Autor (2011). Several papers have stressed the decline in both real wages and weeks worked among men. See Juhn, Murphy, and Topel (1991, 2002) and Moffitt (2012).

[26] One caveat is that states cap the weekly benefit amount. However, combining the ETA's estimates of replacement rates with earnings data from the 2015 American Community Survey shows that at least three-quarters of noncollege educated prime-age men would have received less than the maximum benefit if they were to become unemployed.

[27] ETA's estimates will miss the effect of changes in the duration of benefits on UI-implied replacement rate. However, such changes tend to be temporary, as when Congress extends the duration of benefits in recessions (Rothstein 2011 and Hagedorn et al 2016).

[28] According to Autor and Duggan (2003, 2006), the increase in SSDI replacement rates is likely considerably higher after accounting for the availability of Medicare under SSDI.

benefits are (roughly) indexed to inflation, not past earnings.[29] In light of the fall in real wages, these aspects of SSDI and SNAP would seem to suggest a decline in the income gradient.

However, increases in certain programs' replacement rates do not necessarily translate into a substantial increase in the *average* replacement rate. Consider, again, SSDI. [30] Over our sample period, just 2.5 percent of prime-age men, on average, drew SSDI benefits in a given year. Changes in benefit income among this sample is unlikely to drive noticeable changes in the average income gradient. Moreover, whereas the share of older men (age 55-64) who receive federal disability benefits increased by nearly 8.5 percentage points, [31] the share of prime-age men who receive benefits has risen by only 2 percentage points. Thus, changes in SSDI policy are also unlikely to yield substantial movements in the gradient, or in nonemployment. Indeed, the increase in SSDI receipt can account for at most one-third of the rise in full-year nonemployment, and probably much less than that.[32]

Similarly, the quantitative significance of SNAP/Food Stamps is limited. Between 1979 and 2007, participation of households in our sample hardly budged on net, hovering around just 8 percent.[33] The only meaningful increase in SNAP participation has occurred during the last 10 years, with the participation rate rising to 16 percent by 2015.[34] Even so, SNAP is a relatively small source of nonwork income, even among households whose primary male is out of work for the entire year. Among prime-age males who work less than 13 weeks of the year, higher SNAP participation would imply an increase in average annual household income of only about 3.6 percent since 2007.[35]

Another reason one may expect to see a stronger decline in the income gradient centers around the rise in female labor force participation. The entry of women into the workforce in the last 5 decades has increased the spousal income available to married men. Thus,

---

[29] The SNAP benefit starts at a maximum level and falls at a rate of 30 percent of weekly earnings. The maximum benefit, which is indexed, more specifically, to food prices, rose in real terms by 2.5 percent during our sample.

[30] Data are from the Social Security Administration and show the number of male beneficiaries under 54. There are virtually no recipients under age 25. Note that these calculations rely on records that do not specify beneficiaries' education level, so they refer to all prime-age men.

[31] The Social Security Amendments of 1984 broadened the medical criteria used in the application process.

[32] von Wachter et al (2011) finds that around one-half of rejected SSDI applicants—40 percent of men aged 45-64 and 60 percent of men aged 30-44—return to work. Noting that these estimates are very likely an upper bound on employment rates of SSDI enrollees (Bound, 1989), SSDI can probably account for no more than one-fifth ($0.6 \times 1/3$) of the rise in full-year nonemployment.

[33] Data from the Department of Agriculture. Using the Trim-simulated Food Stamp/SNAP receipt yields a similar picture during the 1990s and 2000s: receipt of income was roughly steady at 13 percent from the mid-1990s through the mid-2000s when it surged to nearly 30 percent during the last recession.

[34] Following Department of Agriculture guidance, roughly 40 states by 2011 had expanded eligibility by lifting income and asset limits on SNAP applicants. See Ganong and Liebman (2013) and Mulligan (2012).

[35] This is calculated by comparing average food stamp benefits per household with what would have been observed assuming the 2007 participation rate prevailed in future years (with everything else being equal).

declines in male labor supply in recent years are more likely to be offset by rising female incomes, reducing the income gradient.

In fact, though, low-skilled males are not more likely to be living with partners, and in particular ones that work. Even after accounting for the rise in cohabitation, the share of noncollege-educated primary men living with a spouse or female partner dropped from 93 percent in the late 1960s to 70 percent today. As a result, despite the increase in female labor force participation, the share of primary males living with a partner that earned labor income was roughly unchanged at 45 percent from the late 1960s to 2015. In fact, among *out-of-work* men in this group, the share with a partner that worked *fell* from 40 percent to 30 percent over that time period (see also CEA, 2016 and Murphy and Topel, 2001). These facts underlie the result in Figure 2, which shows that spousal earnings increased as a share of household income, but not dramatically.

3. Lessons from labor supply theory

In this section we explore the extent to which a coherent account of the rise in male joblessness can be provided by labor supply theory. We show that standard specifications of preferences emphasize movements in the distribution of *replacement rates*—the ratio of nonwork benefits to wages—as key drivers of long-run changes in nonemployment. We further show that, for plausible specifications of individual heterogeneity and selection, any such changes in replacement rates are in turn predicted to leave a clear imprint on the joint distribution of income and employment that we documented in section 2: Higher replacement rates flatten the reduced-form relationship between income and weeks worked. Quantitative predictions from labor supply theory are then confronted with the data presented in the preceding sections.

3.1 A benchmark model

We begin by examining the implications of a simple model of labor supply. Its structure is informed by the particular question we seek to address, namely movements in nonemployment that have evolved and persisted over several decades. The essential trade-offs that inform such long-run changes in employment are captured in large part by a static model of labor supply. Clearly, such a model must necessarily abstract from the timing of individual labor supply responses within the lifecycle, uncertainty in the evolution of income sources, and their potential interaction with an individual's ability to adjust his assets.

Consider an individual i who has preferences over consumption c_i and labor supply $h_i \in [0,1]$ given by

$$\ln c_i - \gamma_i \nu(h_i). \tag{2}$$

Two features of these preferences are worth noting. First, the logarithmic specification of consumption preferences implies a balanced-growth property whereby the income and substitution effects of wage changes cancel. This has become a standard benchmark in the literature on labor supply, so we shall begin with these preferences, and examine deviations from them in future work. A second feature of the preferences in (2) is that the disutility of labor supply $\gamma_i \nu(h_i)$ allows for a dimension of individual heterogeneity: γ_i reflects individual i's idiosyncratic preference for leisure relative to consumption.

Each individual seeks to maximize his utility subject to a budget set comprised of three sources of income, each of which is motivated by the empirical sources documented in section 2. The first is *wage income* that is proportional to the fraction of the year worked, $w_i h_i$, and corresponds to the wage and salary income in Figure 2. The second, *replacement income*, tapers off with the fraction of the year worked, $b_i(1-h_i)$. This plays the role of public benefits—Unemployment Insurance, Social Security, Public Assistance and so on—documented in Figure 2.[36] The third source of income, which we shall refer to as *nonlabor income*, denoted z_i, is taken to be independent of weeks worked at the individual level—that is, individual i receives this income regardless of how much he individually works. We use this to capture the contribution to household income of spouses, cohabiting partners, and other household members in Figure 2.[37]

Note that wages w_i, nonwork benefits b_i and nonlabor income z_i all are allowed to vary across individuals, adding further dimensions of heterogeneity. Thus, the budget constraint faced by individual i is given by

$$c_i \leq w_i h_i + b_i(1-h_i) + z_i \equiv y_i. \tag{3}$$

A consequence of balanced-growth preferences is that, for any given leisure preference γ_i, optimal labor supply in this environment is determined uniquely by the *replacement rates* an individual faces—that is, the ratios of nonwork benefits and nonlabor income to in-work wages, $\beta_i \equiv b_i/w_i$ and $\zeta_i \equiv z_i/w_i$. Interior choices $h_i^* \in (0,1)$ that equate the marginal disutility of labor supply to its marginal benefit imply a fraction of the year worked defined implicitly by

$$\gamma_i \nu'(h_i^*) = \left(h_i^* + \frac{\beta_i + \zeta_i}{1-\beta_i}\right)^{-1}. \tag{4}$$

Corner solutions for full-year nonemployment ($h_i^* = 0$) and full-year employment ($h_i^* = 1$) are given by

[36] The advantage of this specification of benefit income is simplicity. As such, it is, in our view, the right first step to take in assessing the role of replacement rate changes in shaping the long-run path of employment. In ongoing work, we model more explicitly the eligibility criteria of the benefit programs.

[37] We show in the Appendix that a simple unitary model of household labor supply in which household leisure preferences are separable in household members' weeks worked is consistent with this interpretation.

$$h_i^* = 0 \text{ if } \frac{\beta_i + \zeta_i}{1 - \beta_i} > \frac{1}{\gamma_i \nu'(0)}, \text{ and } h_i^* = 1 \text{ if } \frac{\beta_i + \zeta_i}{1 - \beta_i} < \frac{1}{\gamma_i \nu'(1)} - 1. \tag{5}$$

According to equations (4) and (5), optimal labor supply is thus (weakly) decreasing in both the replacement rates β_i and ζ_i, as well as leisure preference γ_i.

Despite the parsimony of the model's structure, the heterogeneity along replacement rates and preferences is fairly rich. Indeed, as we shall see, such heterogeneity can rationalize the cross-sectional differences in weeks worked and incomes observed in Section 2.

By the same token, a key implication is that the large rises in aggregate nonemployment among prime-aged men seen in the data can be accounted for in this framework only by adverse shifts in the distributions of the replacement rates β_i, ζ_i and/or leisure preferences γ_i. Our focus is on whether an explanation based on adverse shifts in replacement rates provides a coherent account of related empirical evidence, and so we hold fixed the distribution of leisure preference in what follows.[38] In this section, we will seek to identify whether movements in average replacement rates in our model can rationalize both the increase in nonemployment as well as a modest decline in reduced-from income gradients that we laid out in Section 2. To see what labor supply theory predicts for the relationship between nonemplyoment and income gradients, it is instructive to consider two polar cases.

A first, simple example is that of degenerate distributions of replacement rates, $\beta_i \equiv \mu_\beta$ and $\zeta_i \equiv \mu_\zeta$ in which case the ratio of wages to benefits is the same for all individuals whether they work or not. The cross section of labor supply in this simplified setting is purely determined by preference heterogeneity. Now, consider a decrease in average wage offers which will in turn induce a increase in aggregate nonemployment. Since there is no heterogeneity with respect to β_i, the workers who cut their labor supply do not, on average, face different wage or benefit draws than workers who continue to work, say, year-around. Hence, the selection mechanism is muted, and the decline in wages simply passes through one-for-one to full-year workers' income. Since average benefits have not changed, the observed relationship between weeks worked and incomes will flatten.

A key question, of course, is whether this prediction is robust to plausible forms of selection. Suppose now there is heterogeneity in the ratio of benefit to wage offers, β_i, and again, consider the case of an increase in mean of β_i, due to a decline in average wage offers. In this case, selection out of employment and into nonemployment is likely to attenuate the effect of this change in the average replacement rate on the reduced-form gradient. The intuition behind this is straightforward: the workers who move out of employment will be those whose replacement rates are high, and wages low, relative to

[38] Some recent work has explored the parallel possibility that leisure has become more valuable over time (see Aguiar, Bils, Charles and Hurst 2017). This work tends to focus more specifically on relatively young men.

the rest of the employed. As a result of this shift in the composition of the workforce, wage income falls by less than the average wage draw, which buoys the income gradient. At the same time, these workers' replacement rates will be low relative to the rest of the *non*employed. Since a low β_i goes hand in hand with a low b_i, average benefit income will tend to fall among low-h workers. The latter also works against a flatter profile of income by weeks.

In somewhat special cases, one can characterize the equilibrium implications of this selection mechanism. For instance, if we neglect spousal income and assume β_i and w_i are jointly log-Normal, then it is possible to show analytically that the gradient must still increase in the average replacement rate, $\mathrm{E}[\beta_i]$.[39]

In general, however, the equilibrium outcome for the income gradient is difficult to characterize. As a result, we pursue the question within a calibrated version of the model of this section.

3.2 Quantitative illustration

This subsection sets up a quantitative variant of the model of section 3.2 to further investigate how aggregate nonemployment and the weeks-income profile react to shifts in mean replacement rates. We calibrate the model to match salient features of the CPS data in the earlier years of our sample on weeks worked and household income for prime-aged primary males with at most a high school degree.

Parameters. With respect to leisure preferences, we have to select a functional form for $v'(h_i)$ as well as a distribution function for the shift term, γ_i. The marginal disutility of labor has a standard isoelastic specification, $v'(h_i) = h_i^{1/\epsilon}$, where ϵ is the Frisch elasticity of labor supply. The distribution of γ_i is semi-parameteric and thus sufficiently flexible to capture the extent of both full-year employment and nonemployment. Specifically, we assume the density function of γ_i is piecewise linear with two "knots" (points where the individual line segments join up). This implies 5 parameters—two knots and three slopes—to pin down. For simplicity (and in the absence of clear direction from the data), we assume γ_i is independent of the replacement rates, β_i and ζ_i, as well as wages.[40]

[39] This proof relies heavily on the tractability offered by the Gaussian distribution, which implies that conditional expectations (e.g., the average wage conditional on a sufficiently low replacement rate) are linear functions. The proof is available on request.

[40] Within a richer theoretical setting, one could perhaps more easily rationalize a correlation between γ_i and the replacement rates. For instance, consider a model of learning by doing, in which a smaller disutility of labor facilitates the accumulation of human capital and leads to higher earnings.

Next, we turn to the distribution of wages, β_i, and ζ_i. We assume the triplet $(\ln w, \ln \beta, \ln \zeta)$ is drawn from a multivariate Normal distribution.[41] Specifically, we assume the benefit income replacement rate is generated by

$$\ln \beta_i = \mu_\beta - \theta_\beta(\ln w_i - \mu_w) + \varepsilon_i,$$

where $\ln w_i \sim N(\mu_w, \sigma_w^2)$ and $\varepsilon_i \sim N(\mu_\varepsilon, \sigma_\varepsilon^2)$. This specification captures the fact that federal law typically specifies benefits as a function of past or present wages. The case of $\theta_\beta \in [0,1]$, in which benefit income b is (at least weakly) positively related to wages, is consistent with the structure of UI and SSDI programs. Means-tested benefits, such as SNAP and TANF, imply instead that benefits fall as wages rise. Our calibration of θ_β will reflect the average relation between benefit and wage draws implied by the benefits system. The relationship between non-labor income and wages is similarly structured

$$\ln \zeta_i = \mu_\zeta - \theta_\zeta(\ln w_i - \mu_w) + \xi_i,$$

with $\xi_i \sim N(\mu_\varepsilon, \sigma_\varepsilon^2)$. The relationship between wages and z can be thought to capture, for instance, the well-known observation that similarly educated men and women are more likely to marry (Mare, 1991; Greenwood et al 2014).

In total, then, we have to calibrate up to 14 parameters: the Frisch elasticity, ϵ; five parameters governing the distribution of preferences, γ; the means and variances of $\ln w, \ln \beta,$ and $\ln \zeta$; and, finally, θ_β and θ_ζ. In practice, though, we only have to pin down 12 of these, and the other two are implied.

Note, first, that our assumption on preferences implies that the mean of $\ln \beta$ can be inferred immediately once its standard deviation, σ_β, is known. To see this, note that the marginal disutility of labor is zero at $h = 0$. Hence, a worker chooses $h = 0$ only if $\beta > 1$. Therefore, the full-year nonemployment rate is given by $\Pr(\beta > 1) = 1 - \Phi\left((\ln 1 - \mu_\beta)/\sigma_\beta\right)$. Thus, for any σ_β, we can choose μ_β to replicate the observed full-year nonemployment rate, which is one of the more salient moments in the data.

Next, we do not attempt to identify σ_ε^2. We found that for plausible values of other parameters, the data would "ask" for (much) more variance in $\ln w$ than in $\ln \beta$. Unless θ_β is ratcheted down near zero, the only way to accommodate this is to make the variance of ε negative—an obvious impossibility. To simplify the calibration, then, we fix $\sigma_\varepsilon^2 = 0$, in which case values for σ_β^2 and θ_β immediately imply the variance of wages.

[41] Equivalently, one could write the triplet in terms of wages, benefits (b), and other household members' income (z). We cast it in terms of the (log) replacement rates in light of their salience within the model.

Moments and Identification. Broadly, we target two sets of moments from data in the early years of our CPS sample. One is the weeks worked distribution. In particular, we target the share of primary males in six weeks worked bins: zero weeks, 1-12 weeks, 13-25 weeks, 26-38 weeks, 39-50 weeks, and 51-52 weeks (full year work). Given the salience of full-year nonemployment and employment, we separate these from their nearest bins. As we turn next to moments of the household income distribution, we will, for the sake of parsimony, typically fold the zero-weeks bin into the one just above it (1-12 weeks) and the full-year bin into the one just below it (39-50 weeks).

We also target the relation between various sources of household income and the primary male's weeks of work. For starters, we target average primary male wage income within each of four weeks bins: 0-12 weeks, 13-25 weeks, 26-38 weeks, and 39+ weeks. The profile of wage income by weeks is remarkably linear (see Figure 2), which conveys valuable information about the structure of heterogeneity. To see this, suppose variation in β were the sole source of heterogeneity. Since a higher wage is manifest in a lower replacement rate, wages and weeks would be strongly positively correlated. This implies a more convex relationship between weeks (h) and wage income (wh), since an increase in weeks will go hand in hand with a higher wage rate (w). In contrast, if γ represented the sole source of heterogeneity, the cross section of weeks worked would be independent of w, b, and z. Therefore, the relation between wage income and weeks would in fact be linear, with a slope of $\mathrm{E}[w]$. This observation previews why heterogeneity in our calibrated model is dominated by variation in γ.

In addition, we target the relation between the primary male's weeks worked and other members' income. There are two parts to this. The first targets average household benefit income within each of the four weeks bins noted above. The second targets average labor income among other household members within each of these four bins. These moments convey considerable information regarding, among others, the slope parameters θ_β and θ_ζ. In the case of benefit income, for instance, the data (Figure 2) indicate that it falls off precipitously with weeks worked (of the primary male). This feature of the data will militate against any strong positive correlation between benefit (b) and wage (w) draws. Otherwise, workers who draw modestly high wages and work part-year will also tend to collect a considerable sum in benefit income, which runs against what we see in the data.

Last, but certainly not least, we target the income gradient. We generate this moment within the model just as we do in the data. Specifically, we project simulated log household income on the share of year worked.

In total, then, we target 19 moments. Six of these pertain to the cross section of weeks worked. The other relate to the profile of household income by weeks worked.

Results. The calibrated model matches data on weeks and income reasonably well in our view. Figure 6 shows the weeks worked distribution in the model and data. The model does understate the share of workers between 1 and 26 weeks and overstate the share between 27 and 50. However, it very nearly replicates the extent of full-year employment (as well as, of course full-year nonemployment).

Consider next the income-related moments. With respect to the profile of wage income by weeks, the model again does reasonably well in targeting the very high (39+) and very low (less than 12) weeks bins (see Table 2). The model's errors are concentrated in the interior of the weeks distribution; it overstates wage income in the 27-39 weeks bin and understates it in the 13-26 weeks bin. The model does a remarkably good job of getting the gradient of benefit income by weeks. In other words, it captures the precipitous decline in benefit income as weeks increase above 12. The profile of other members' income by weeks shows much less of a gradient in both data and model. That said, there is a slight convexity to this relation in the data that the model does not replicate.

Figure 7 summarizes the profile of household income, and its components, by weeks worked (of the primary male). In this context, the most noticeable error is the overstatement of wage income in the 27-39 weeks bin. Still, the goodness of fit in the low (less than 12 weeks) and high (39+) weeks bins bodes well, since most of the population is in these bins. Hence, when we run the income gradient regressions in the model, the result (0.88) reasonably replicates its empirical counterpart (0.80).

Comparative static. We now use the calibrated model to investigate how a change in the average benefit income replacement rate shapes aggregate nonemployment and the reduced-form income gradient. The decline in μ_β is triggered by a fall in the average wage draw, μ_w.[42]

There are a few different ways to adjust the wage. What we do here is identify a reduction in μ_w to replicate the rise in nonemployment that we observe in the CPS. This does call for a counterfactually large decline in average wages (of 37 log points), but allows us to evaluate the income gradient given the observed change in nonemployment. Alternatively, we could have lowered μ_β to generate a decline in line with the data (roughly, 15 log points). The change in the gradient *relative to* the change in nonemployment is roughly the same in either case.

As we anticipated in section 2, the change in μ_w does indeed imply a change in the reduced-form gradient of the same sign, despite the presence of selection. The model-implied income gradient declines by 13.5 percentage points. This exceeds what we see in the data, but not by too much. At one level, then, this result is a victory for the model.

[42] One can rather easily imagine that policymakers hold fixed benefits, b, in the face of a one-time decline in wages. It is arguably more difficult to argue that the marriage market is unaffected by large changes in wages. This is why we do not hold other members' income, z, constant. In fact, we do just the opposite: we assume z changes proportionately with wages, such that the replacement rate out of other members' income is unchanged.

However, what we saw in section 2 is that, in the data, changes in the gradient are only loosely related to changes in nonemployment. This is especially true for the case of *full-year* nonemployment, which trends relentlessly upward despite medium-run fluctuations in the gradient.

To illustrate the extent to which this poses a challenge for the model, we carry out a simple exercise. In the model, the full-year nonemployment rate increases by 8.4 percentage points, which undershoots the 11 percentage point increase in the data. As a result, the gradient in the model is particularly responsive relative to the full-year nonemployment rate: the gradient decreases by 1.6 percentage points per percentage point increase in the full-year rate. The counterpart to this in the data is 0.9. To see the implications of this, we take decadal changes in the gradient from the data and predict the change the full-year rate based on the model's elasticity of 1.6. The application to decadal data acknowledges that ours is a static model and, thus, best suited to address longer-run developments.

Figure 8 presents the results. While the full-year nonemployment rate increased in each subperiod, the model's prediction is considerably out of step. For example, the model predicts an enormous decline in the full-year nonemployment rate between 1977 and 1987 when in fact it increased slightly. In addition, the predicted increase in nonemployment since 2000 is very modest relative to the data. Instead, much of the increase in nonemployment in the model occurs between 1987 and 1997, but the model's prediction exceeds the observed change by a factor of three.

3.3 Further theoretical explorations and summary (in progress)

In light of the results of Section 3, this section will explore several refinements of canonical labor supply theory and assesses their ability to rationalize a rise in nonemployment that is not accompanied by a decline in the income gradient over the same time period. Specifically, we explore the robustness of the result in richer dynamic, stochastic environments, and to deviations from the assumptions underlying the above benchmark model.

Accommodating Disability Insurance. Our model presently allows for a unit taper of benefit income with hours worked. However, some programs are all or nothing: if you work, you cannot receive the benefit. The most prominent such program is SSDI. We will explore extensions of the model to incorporate a program of this structure, and more realistic benefit eligibility criteria.

Higher-order moments. We will experiment with alternative comparative static exercises that adjust the wage offer distribution at the tails. For example, a sharp decline in wage offers mainly for the lower end of the distribution.

Non balanced-growth preferences. We will consider relaxing the restriction on balanced growth preferences such as plausible deviations from balance growth that typically involve elasticities of intertemporal substitution that are *less than* one. This implies income effects of a wage change exceed substitution effects. As a result, the model is likely to predict smaller declines in employment over our sample period for the same decline in wages.

Dynamics with static expectations. In this extension, the agent is embedded into a dynamic problem in which he chooses a sequence of consumption and labor supply to maximize discounted lifetime utility. The individual can borrow and save at a risk-free rate, subject only to a no-Ponzi condition. This is potentially a rich problem, but the assumption of static expectations simplifies it dramatically. Under static expectations, the individual expects, in each period, that all future benefit and wage rates will equal their current values. Since any future period will look like the present with probability one, this model collapses to the static decision problem already analyzed.

Dynamics with rational expectations. Consider the same dynamic environment just outlined, but now imagine the agent is endowed with rational expectations over future benefit and wage rates. This alteration is unlikely to substantially affect our conclusions. In reality, public assistance programs tie benefit growth to wages, and individual wage processes are very persistent. As a result, the solution under static expectations, though not quite optimal, is likely to provide a good approximation to the properties of dynamic models such as this one.[43]

Borrowing constraints. Continuing on in the dynamic setting, now consider the effects of introducing borrowing constraints. This addition to the model is, again, unlikely to alter dramatically the implications of the static theory. Borrowing constraints bind in situations where workers cannot fully anticipate, and plan for, changes in the state. Suppose, then, that workers are "surprised" by a decline in average wages, causing the borrowing constraint to bind. In such a case, the worker should work *more* than otherwise to prevent substantial declines in consumption. Thus, even if we grant that workers were surprised by the downward trend in real wages, this model is likely to predict smaller declines in employment than in the static model.

Hand-to-mouth consumers. The preceding logic carries over to other cases in which, for any number of reasons, workers are unable to smooth consumption as in a canonical life cycle setting. For instance, suppose workers are hand-to-mouth (Kaplan, Violante,

[43] Agents under rational expectations may foresee future declines in aggregate real wages of the sort observed in our sample. However, under balanced growth preferences, these anticipated declines will have offsetting income and substitution effects on labor supply.

Weidner, 2014). This again means that consumption is more responsive to any anticipated decline in earnings, leading to smaller employment responses.

References

Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galgis-Reis, Christopher Smith, and William Wascher. 2014. "Labor Force Participation: Recent Developments and Future Prospects." Brookings Papers on Economic Activity, Fall 2014.

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In David Card and Orley Ashenfelter (eds.) *Handbook of Labor Economics* Volume 4 Part B, Elsevier: 1043-1171.

Aguiar, Mark, Mark Bils, Kerwin Charles, and Erik Hurst. 2017. "Leisure Luxuries and the Labor Supply of Young Men." NBER Working Paper No. 23552.

Autor, David H. and Mark G. Duggan. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment." Quarterly Journal of Economics, February 2003, 118(1), 157-205.

Autor, David H. and Mark G. Duggan. 2006. "The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding." Journal of Economic Perspectives, Summer 2006, 20(3), 71-96.

Blinder, Alan. 1985. "Measuring Income: What Kind Should be In?: Comment" Proceedings of the Conference on the Measurement of Noncash Benefits.

Bound, John. 1989. "The Health and Earnings of Rejected Disability Insurance Applicants." American Economic Review, 79:482-503.

Casper, Lynne M. and Phillip N. Cohen. 2000. "How does POSSLQ measure up? Historical estimates of cohabitation." *Demography*, Volume 37, Issue 2, pp 237-245.

Chetty, Raj. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." *Econometrica* 80(3): 969-1018.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review Papers and Proceedings* 101: 471-75.

Council of Economic Advisors. 2016. "The Long-term Decline in Prime-age Male Labor Force Participation."

Ellwood, David T and Lawrence H. Summers. 1985. "Measuring Income: What Kind Should be In?" Proceedings of the Conference on the Measurement of Noncash Benefits.

Elsby, Michael W. L., and Matthew D. Shapiro. 2012. "Why Does Trend Growth Affect Equilibrium Employment? A New Explanation of an Old Puzzle." *American Economic Review* 102(4): 1378-1413.

Finkelstein, Amy, Neale Mahoney, and Matthew J. Notowidigdo. 2018. "What Does (Formal) Health Insurance Do, and For Whom?" Forthcoming, *Annual Review of Economics.*

Fitch, Catherine, Ron Goeken, and Steven Ruggles. 2005. "The Rise of Cohabitation in the United States: New Historical Estimates." Mimeo, Minnesota Population Center, University of Minnesota.

Flood, Sarah, Miriam King, Steven Ruggles, and J. Robert Warren. Integrated Public Use Microdata Series, Current Population Survey: Version 5.0. [dataset]. Minneapolis: University of Minnesota, 2017. https://doi.org/10.18128/D030.V5.0.

Ganong, Peter and Jeffrey B. Liebman. 2013. "The Decline, Rebound, and Further Rise in SNAP Enrollment: Disentangling Business Cycle Fluctuations and Policy changes." NBER Working Paper No. 19363.

Giannarelli, Linda. 1992. *An Analyst's Guide to TRIM2: The Transfer Income Model Version 2.* Washington, D.C.: The Urban Institute Press.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos. 2014. "Marry Your Like: Assortative Mating and Income Inequality." *American Economic Review P&P*, v. 104, n. 5, 348-353.

Gruber, Jonathan. 2000. "Tax Subsidies for Health Insurance: Evaluating the Costs and Benefits." NBER Working Paper No. 7553.

Hagedorn, Marcus, Fatih Karahan, Iourii Manovski, and Kurt Mitman. 2016. "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects." NBER Working Paper No. 19499.

Janicki, Hubert, Brett O'Hara, and Alice M. Zawacki. 2013. "Comparing Methods for Imputing Employer Health Insurance Contributions in the Current Population Survey." US Census Bureau Center for Economic Studies Paper No. CES-WP- 13-41.

Juhn, Chinhui, Kevin M. Murphy, and Robert Topel. 1991. "Why Has the Natural Rate of Unemployment Increased over Time?" *Brookings Papers on Economic Activity* 1991(1): 75-126.

Juhn, Chinhui, Kevin M. Murphy, and Robert Topel. 2002. "Current Unemployment, Historically Contemplated." *Brookings Papers on Economic Activity* 2002(1): 79-116.

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealth Hand-to-Mouth." Brookings Papers on Economic Activity, Spring.

Keane, Michael, and Richard Rogerson. 2012. "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom." *Journal of Economic Literature* 50(2): 464–476.

Krueger, Alan B. 2017. "Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate." Brooking Papers on Economic Activity Conference Drafts, September.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Aysegul Sahin. 2012. “Gross worker flows over the business cycle.” Working Paper No. 17779, National Bureau of Economic Research.

Mare, Robert D. 1991. “Five Decades of Educational Assortative Mating.” *American Sociological Review*, Vol. 56, No. 1, pp. 15-32.

Meyer, Bruce D., Wallace K. C. Mok, and James X. Sullivan. “Household Surveys in Crisis.” *Journal of Economic Perspectives* 29(4): 199-226.

Moffitt, Robert A. 2012. “The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations.” *Brookings Papers on Economic Activity*.

Muller, L. Scott. 2008. “The Effects of Wage Indexing on Social Security Disability Benefits.” *Social Security Bulletin*, Vol. 68, No. 3.

Mulligan, Casey. 2012. *The Redistribution Recession: How Labor Market Distortions Contracted the Economy*. Oxford University Press.

Organization for Economic Cooperation and Development. 2011. “Divided We Stand—Why Inequality Keeps Rising.”

Peterson, Chris L., and Christine Devere. 2002. “Health Insurance: Federal Data Sources for Analyses of the Uninsured.” Congressional Research Service, Report to Congress.

Purcell, Patrick J. 2015 “Income Taxes on Social Security Benefits.” Social Security Administration Issue Paper 2015-02.

Rees Albert. 1985. “Measuring Income: What Kind Should be In?: Comment” Proceedings of the Conference on the Measurement of Noncash Benefits.

Rothstein, Jesse. 2011. “Unemployment Insurance and Job Search in the Great Recession.” *Brookings Papers on Economic Activity* Fall 2011: 143-210.

Tax Policy Center. 2015. “Average and Marginal Combined Federal Income and Social Security and Medicare (FICA) Employee Tax Rates for Four-Person Families at the Same Relative Positions in the Income Distribution, 1955-2014.” Data report.

Von Wachter, Till, Jae Song, and Joyce Manchester. 2011. “Trends in Employment and Earnings of Allowed and Rejected Applicants to the Social Security Disability Insurance Program.” *American Economic Review*, Vol. 101, No. 7, pp 3308-29.

Zedlewski, Sheila R., and Linda Giannarelli. 2015. “TRIM: A Tool for Social Policy Analysis.” Research Report, Urban Institute.

Figures

Figure 1. Prime-age male nonemployment rates, 1967-2015.

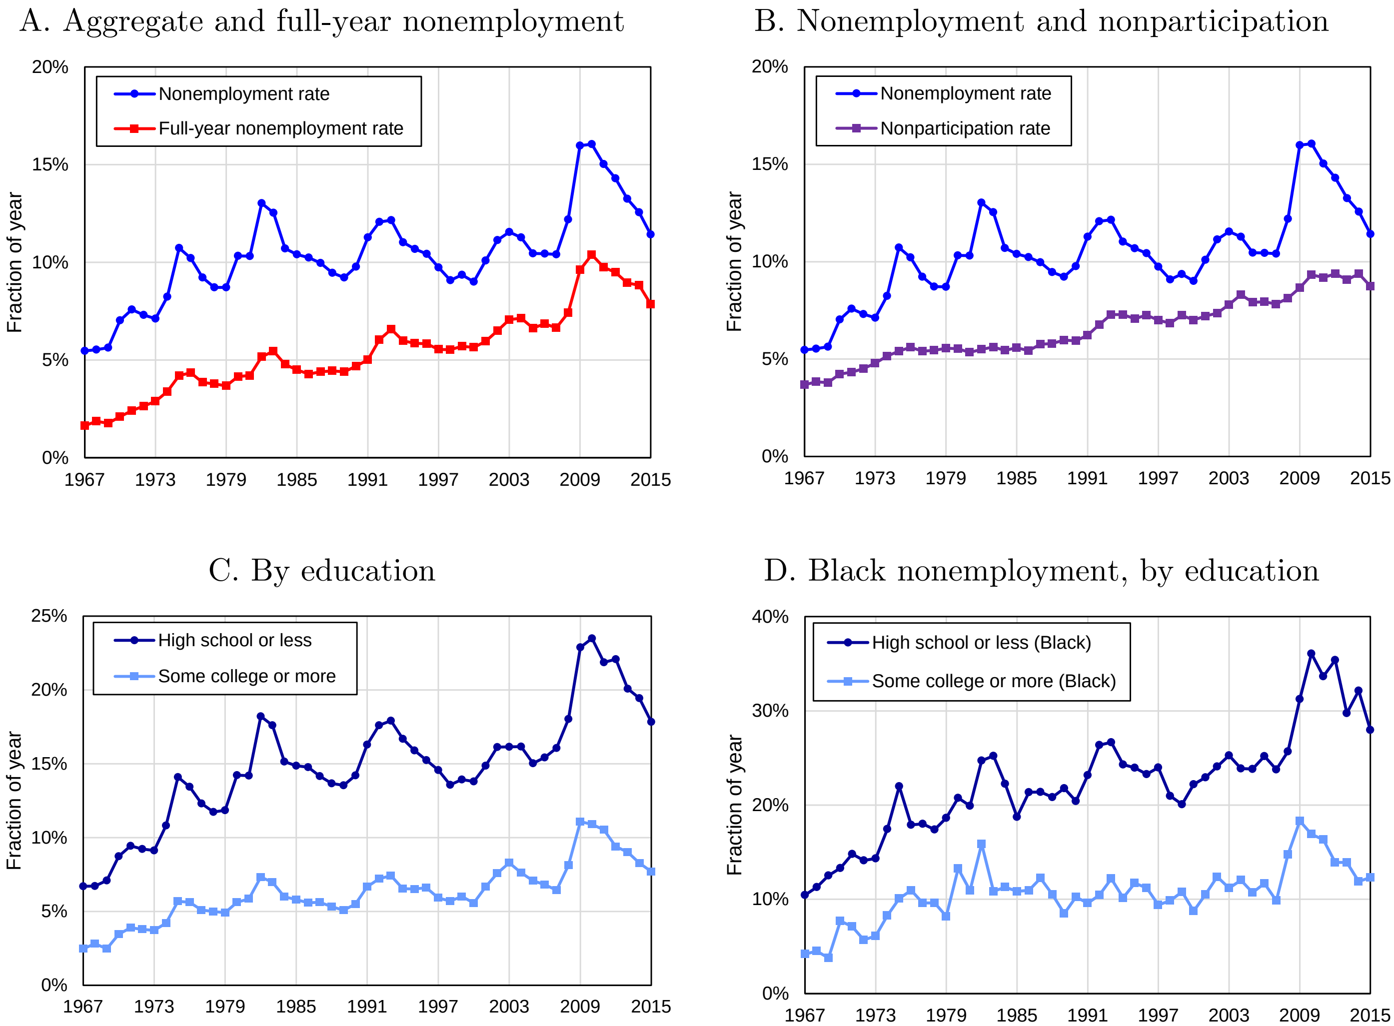

Table 1. Decomposition of income sources in the March Current Population Survey.

Household income
Head's income
Earned income
Wage and salary income
Unearned income
Public assistance
AFDC; TANF; SSI; SNAP/Food Stamps (1979 onwards); other aid from welfare offices
Unemployment
UI benefits; worker's comp.; veteran's; gov't. pensions; non-OASDI retirement/disability
Social Security
All income from Old Age, Survivor and Disability Insurance (OASDI)
Other income
Interest, dividends, rentals; alimony; child support; friends, educational assistance
Spousal/cohabiting partner income
Other household member income

Figure 2. Household income by weeks worked in the March CPS, High school or less, Pre- versus post-1990.

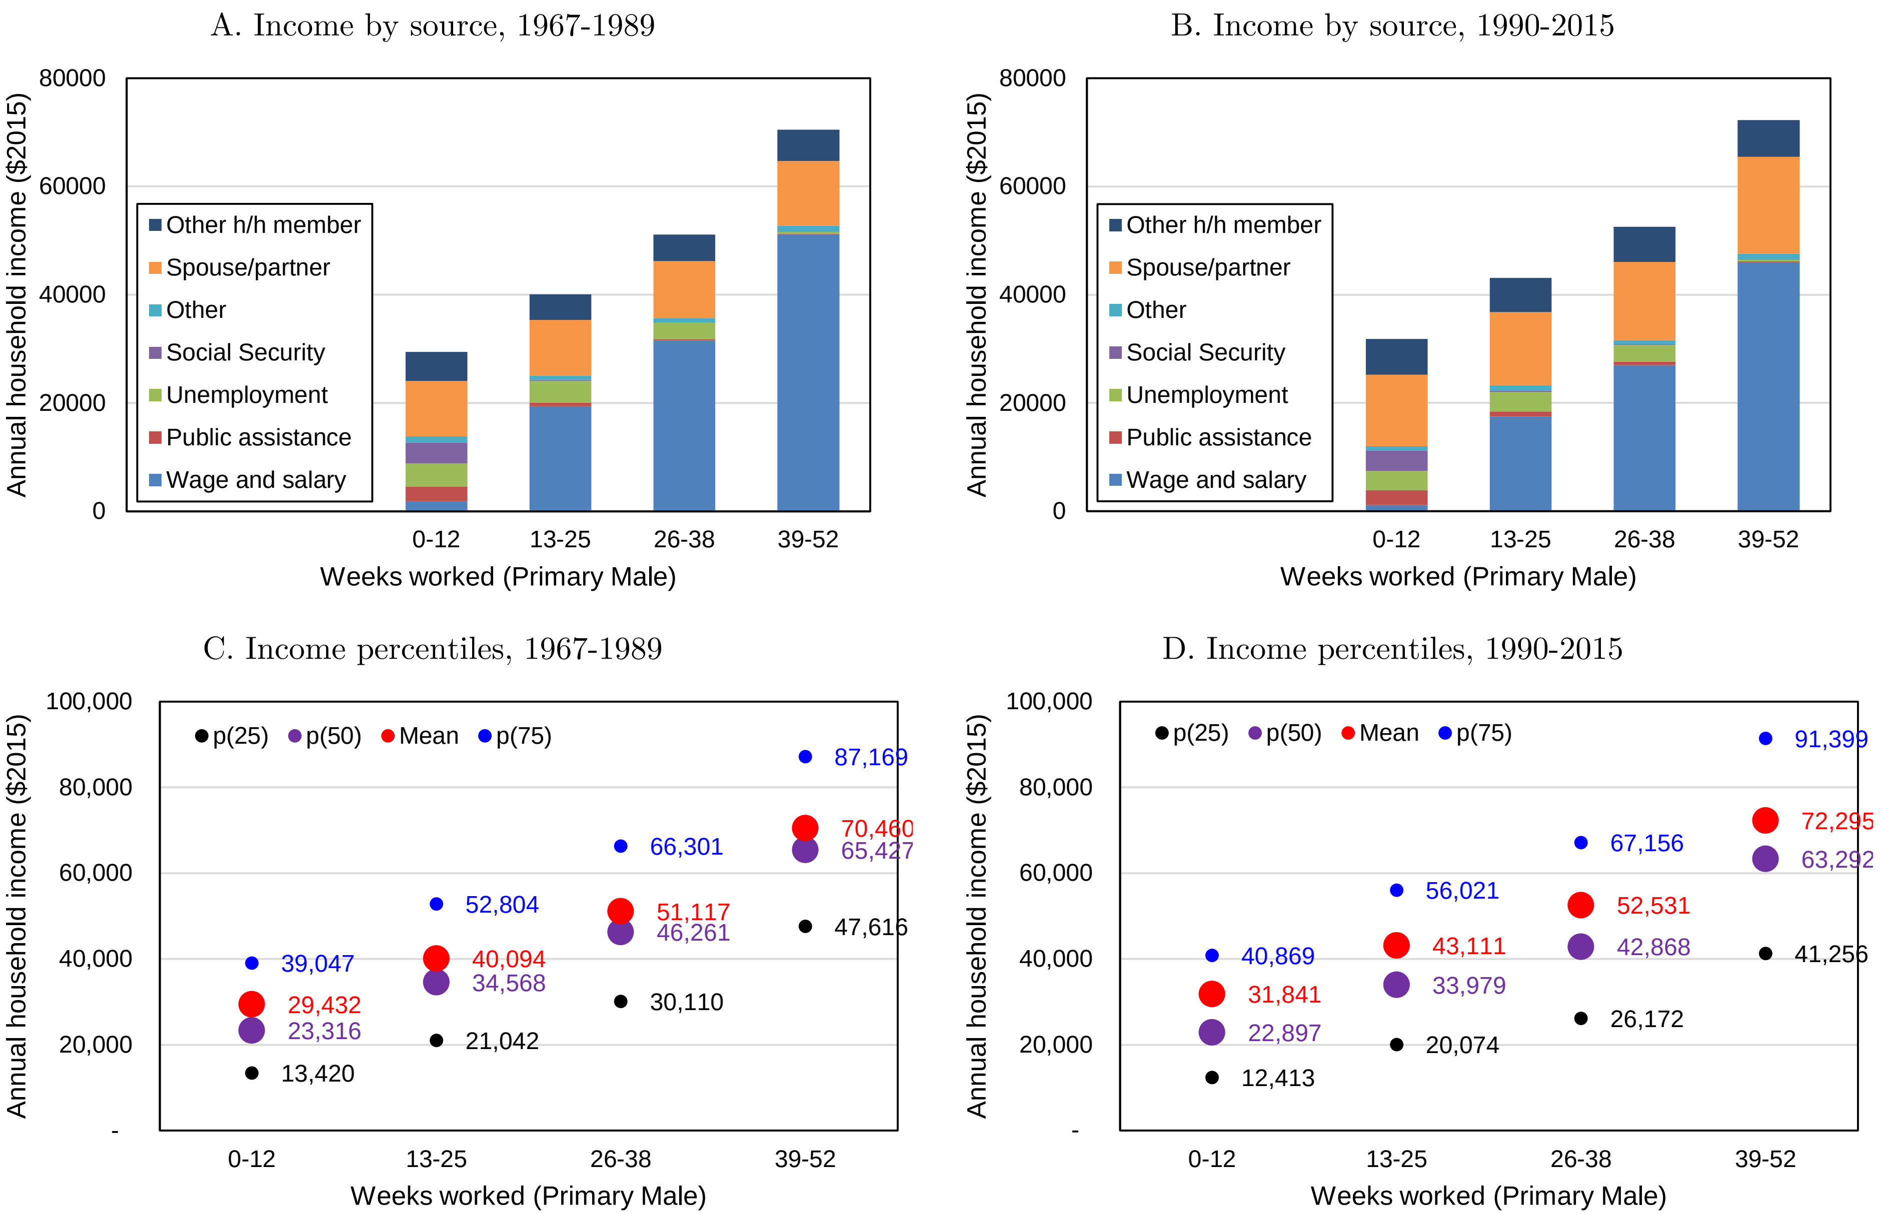

Figure 3. Estimated income gradient over time from the March CPS.

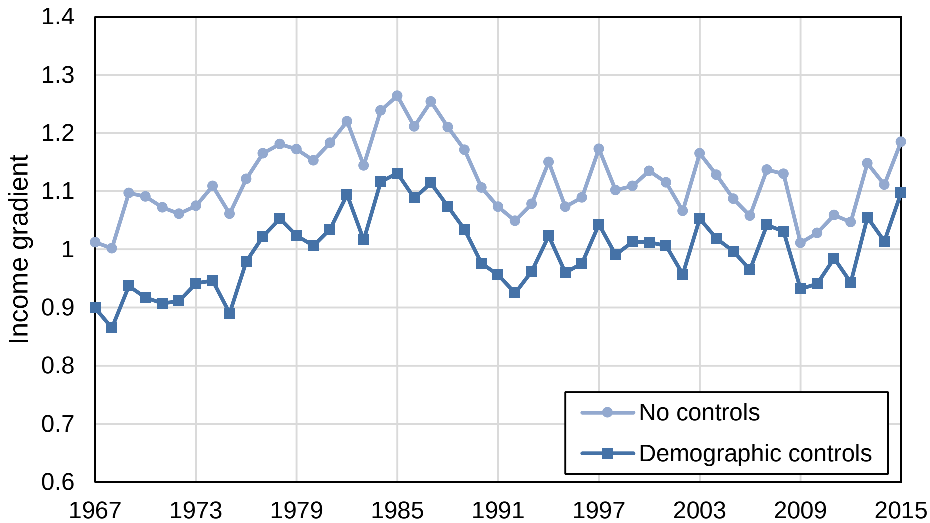

Note: Non-college education primary males. Controls include a quartic in age, quartic in potential labor market experience, number of household members, number of children, race, and an indicator for whether the primary male is a spouse or cohabiting partner of the household head.

Figure 4. Before- and after-tax household income.

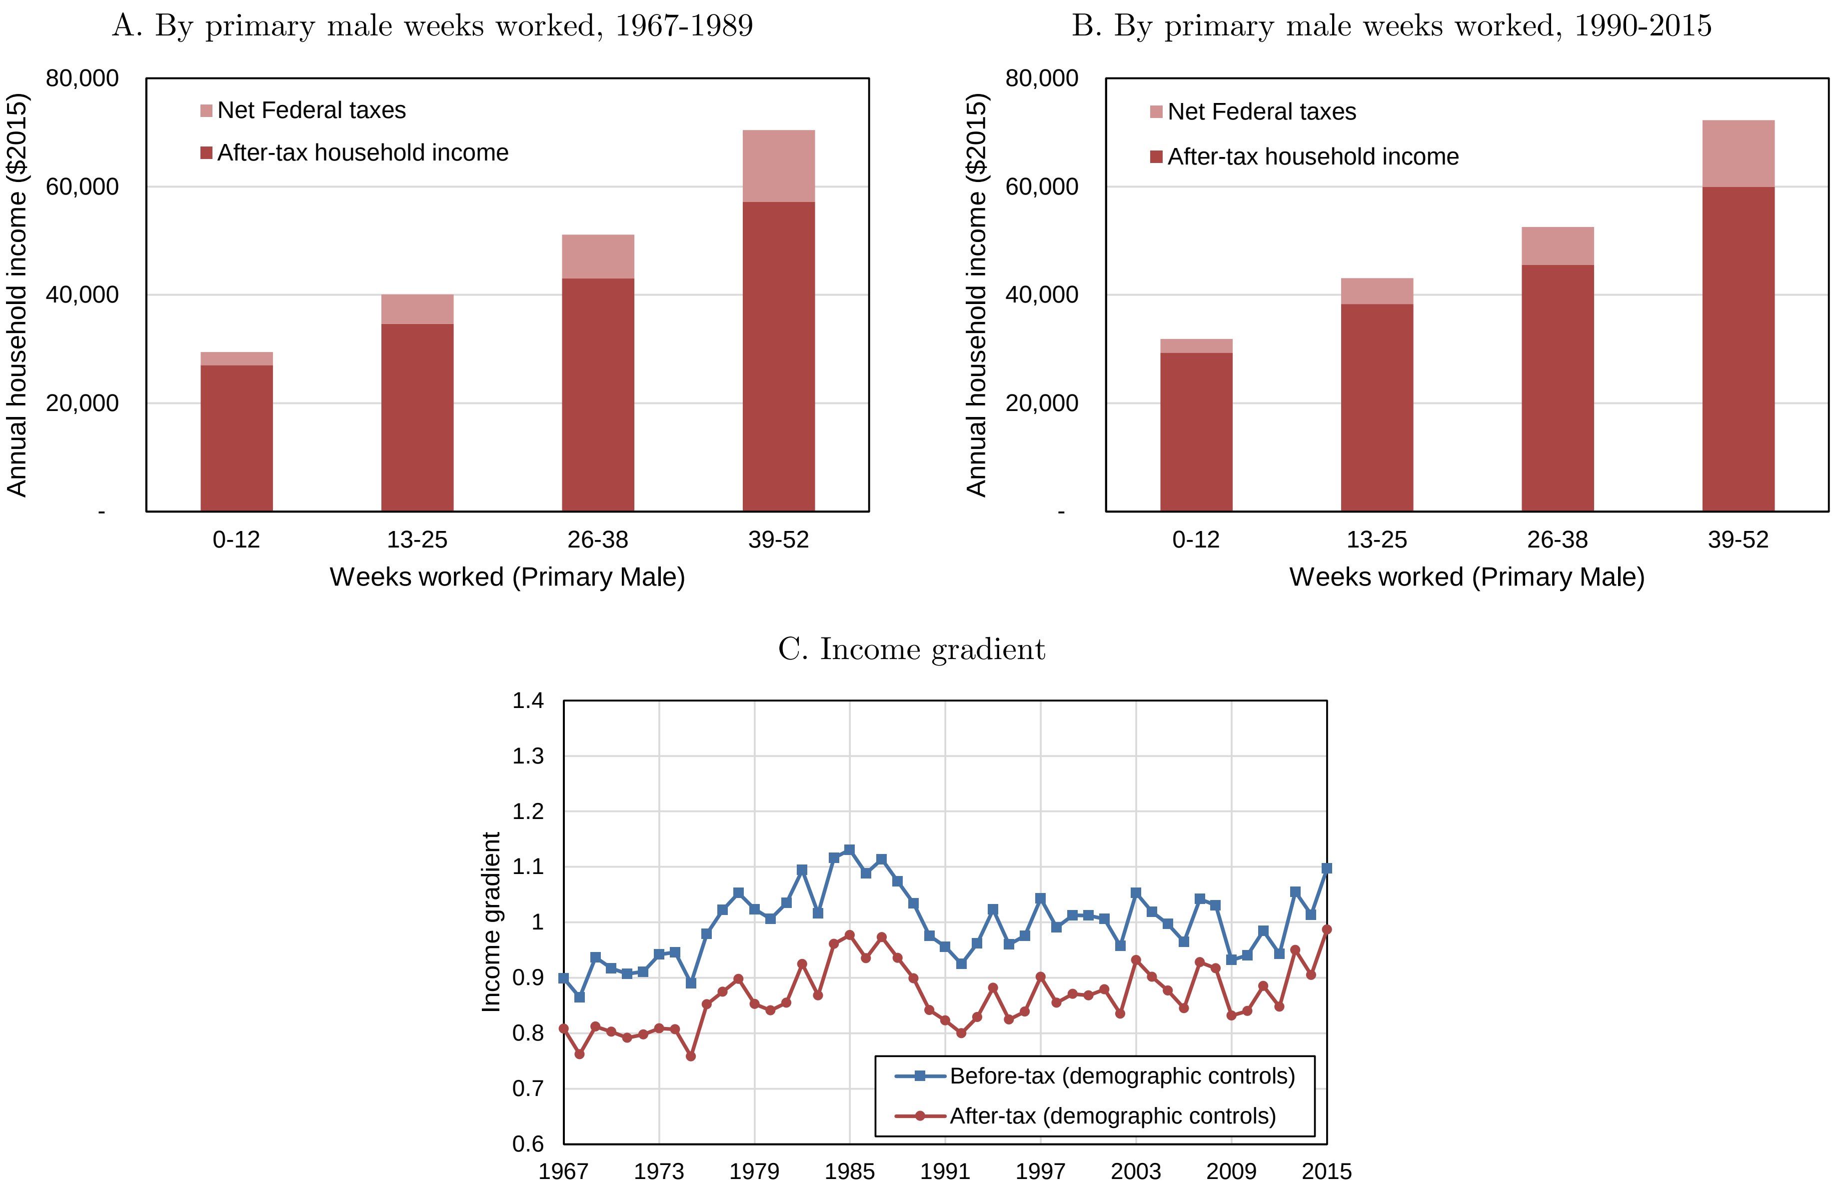

Note: After-tax income is derived using NBER's TAXSIM and calculated as total federal income tax liability plus the employee-share of FICA liability. Before-tax income and controls are the same as in Figure 3.

Figure 5. Estimated income gradient over time, extensions to March CPS.

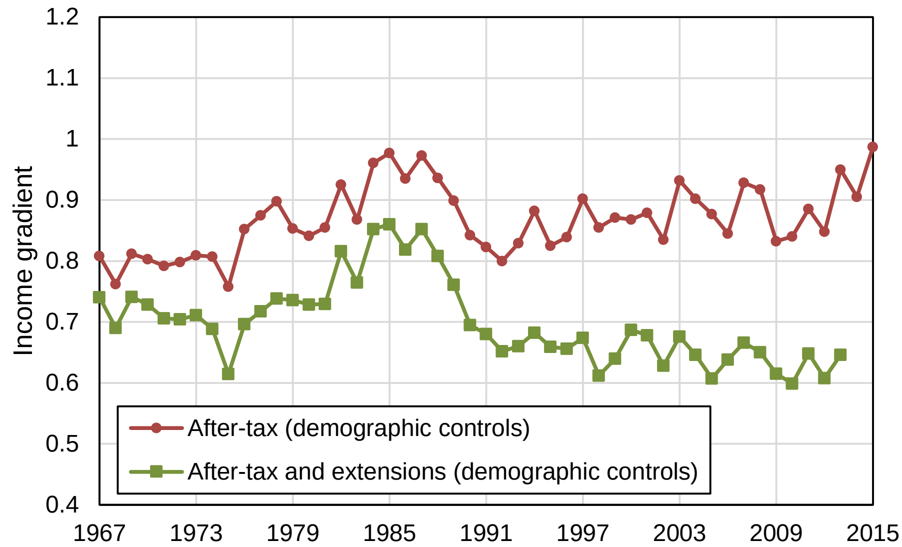

Figure 6. Weeks worked distribution, model and data

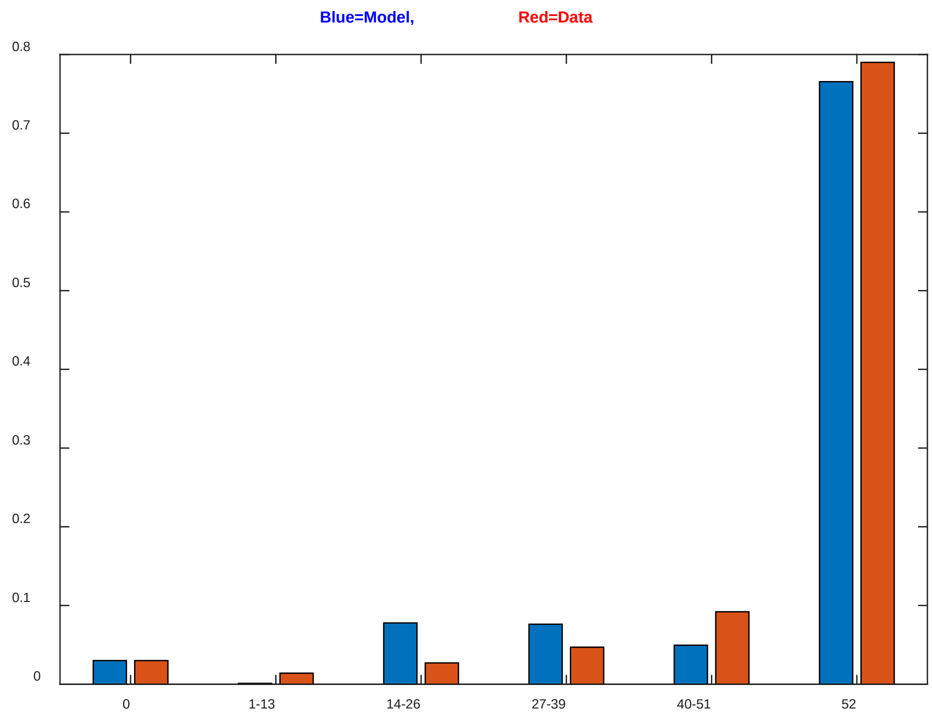

Figure 7. Household income, model and data (thousands of 2015 $)

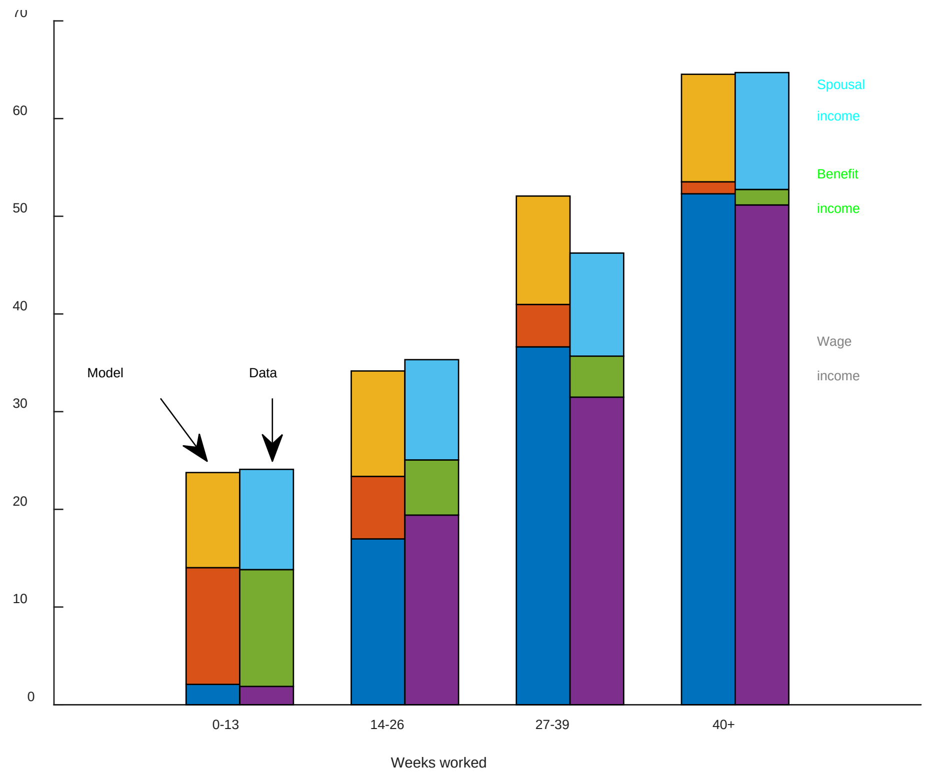

Figure 8. Model prediction for full-year nonemployment versus data

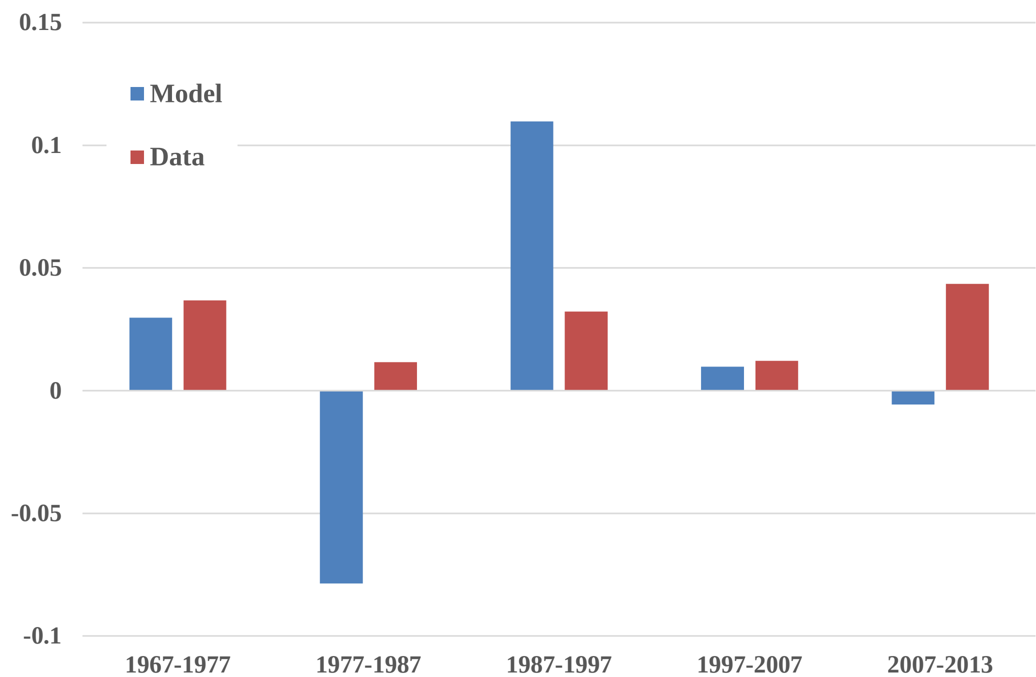

Table 2: Moments in Calibrated Model versus Data

	Model	Data
Pr(h = 0)	0.03	0.03
Pr(h≤.25)	0.03	0.04
Pr(h≤50)	0.11	0.07
Pr(h≤.75)	0.19	0.12
Pr(h=1)	0.77	0.79
E[h]	0.88	0.91
Wage income | h≥.75	$52,302	$51,160
Wage income | h ∈ (.5,.75)	$36,633	$31,484
Wage income | h ∈ (.25,.5)	$16,971	$19,411
Wage income | h ≤.25	$2,081	$1,871
Benefit income | h≥.75	$1,226	$1,587
Benefit income | h ∈ (.5,.75)	$4,341	$4,198
Benefit income | h ∈ (.25,.5)	$6,400	$5,645
Benefit income | h ≤.25	$11,950	$11,959
Non-labor income | h≥.75	$11,006	$11,966
Non-labor income | h ∈ (.5,.75)	$11,096	$10,552
Non-labor income | h ∈ (.25,.5)	$10,793	$10,265
Non-labor income | h ≤.25	$9,734	$10,264
Income gradient	0.88	0.80

Table 3: Model Parameters

ϵ	0.21947
Preferences- slope #1	12.92404
Preferences- slope #2	0.00030
Preferences- slope #3	0.00013
Preferences- knot #1	0.35518
Preferences- knot #2	0.77629
σ^{β}	0.69450
μ_w	6.67020
θ_{β}	1.06482
μ_{ζ}	-1.32916
θ_{ζ}	0.92140
$\sigma_{\epsilon}^{\zeta}$	0.12849

Appendix

A. A simple model of household labor supply

Suppose there is a household with preferences over household consumption c, and male and female weeks worked, h_m and h_f, given by $\ln c - \nu_m(h_m) - \nu_f(h_f)$. The household seeks to maximize the latter subject to the household budget constraint,

$$c \leq w_m h_m + b_m(1 - h_m) + w_f h_f + b_f(1 - h_f). \tag{6}$$

Household income arises from male and female wage income, $w_m h_m$ and $w_f h_f$, as well as male and female benefit income, $b_m(1 - h_m)$ and $b_f(1 - h_f)$. Optimal labor supply choices, $\{h_m^*, h_f^*\} \in [0,1]^2$, will satisfy the first-order conditions

$$\begin{aligned} \nu_m'(h_m^*) &\gtreqless \frac{w_m - b_m}{w_m h_m^* + b_m(1 - h_m^*) + w_f h_f^* + b_f(1 - h_f^*)}, \text{ as } h_m^* \in 0, (0,1), 1, \\ \nu_f'(h_f^*) &\gtreqless \frac{w_f - b_f}{w_m h_m^* + b_m(1 - h_m^*) + w_f h_f^* + b_f(1 - h_f^*)}, \text{ as } h_f^* \in 0, (0,1), 1. \end{aligned} \tag{7}$$

A sufficient statistic for the effect of the female's labor supply on the male's labor supply is thus the income generated by the female, $z \equiv w_f h_f + b_f(1 - h_f)$ (and *vice versa*). The solution to this system will imply that z is a function of (ratios of) the parameters $\{w_m, b_m, w_f, b_f\}$ (given balanced-growth preferences). In this way, model outcomes are as if the male receives nonlabor income z that is a random variable correlated with the parameters of the male's income, $\{w_m, b_m\}$.

B. Further details of adjustments to CPS measures (in progress)

After-tax income. To estimate net federal tax liability (taxes paid minus credits, including the EITC) we use the NBER TAXSIM model. TAXSIM is a set of programs that estimate tax liability from household survey data, and is a widely-used tax simulator. It offers estimates of federal tax liability from 1960 on, and state tax liability from 1977 on. We estimate only federal liability because that allows coverage of our full sample.

TAXSIM requires inputs at the tax-filer level (for example, one tax file is a married couple), whereas the CPS does not have (consistent) data for tax units. Therefore, we must use some rules to determine tax units in the CPS. We largely follow Cooper, Lutz, and Palumbo (2011) to assign individuals to tax units. Individuals over 18 are defined as their own tax filing unit, even if they live with their parents or relatives. Children 15 and over who have positive earnings are also assigned their own tax unit. We classify joint tax units if the household head is married, and sum income across the spouses; if the household head is single without kids, they are sole filers; if they have dependents, they are head of household filers.

There are a number of income categories that can alter tax liability that the CPS does not record (although these are not usually seen as quantitatively important). These include property income, rent paid, child care expenditures, and capital gains, among some other categories. Notably, TAXSIM does estimate the taxable portion of income from transfer programs such as Unemployment Insurance and Social Security. Further, we leave itemized deductions blank and assume that all households take the standard deduction. Household after-tax income is defined as pre-tax income minus federal income tax liability net of credits (including the EITC) minus one-half of FICA tax liability (reflecting the share paid by the household rather than the employer).

Medicaid and Medicare. Here, we report the sources of our data on Medicaid and Medicare expenditure and our construction of state-level estimates.

Medicaid. Our state-level data on Medicaid expenditure per person served are taken from the following three sources.

(i) 1986—1998: Medicaid State Data Tables, published by the Health Care Financing Administration (HCFA).
(ii) 1999—2011: Medicare and Medicaid Statistical Supplement, published by the Centers for Medicare and Medicaid Services (CMS), the successor to the HCFA.
(iii) 2012—2013: Medicaid and CHIP Data Book, published by Medicaid and CHIP Payment and Access Commission (MACPAC).

For each state, expenditure is reported for four classes of recipients: the disabled; children (under 21); adults (age 21-64); and the elderly (over 64).

These data are derived from HCFA Form 2082, which is completed by the state governments. However, Form 2082 data are not the basis of federal Medicaid reimbursements; the latter are based on HCFA Form 64. Accordingly, the Form 64 data are likely to be a more accurate recording of actual spending. Unfortunately, Form 64 does not include enrollment data, and the spending data are not classified by age or disabled status. Some scholars have made occasional attempts to integrate the Form 64 and 2082 files, but these integrated data are unavailable to us.[44] We thus rely on the published tabulations by HCFA, which are based on Form 2082.

The Census relies on the same Form 2082 data to calculate the "market value" of Medicaid. Not surprisingly, our estimates are almost always agree with those carried out by the Census. We deviate in one respect, though. Prior to 1997, participation in AFDC automatically qualified a household for Medicaid. As a result, the HCFA tables reported receipt and expenditure for AFDC children and adults, specifically. Other nondisabled

[44] For instance, see Winterbottom et al. (1995).

and nonelderly beneficiaries were classified as “other Title XIX” recipients, and expenditure for such recipients was reported separately.[45] The Census does not use the other Title XIX data. However, this class of recipients grows steadily following legislation in the mid-to-late 1980s broadening Medicaid eligibility beyond AFDC receipt. We would like to capture these relatively new enrollees. Though we do not know the ages of the beneficiaries under the “other Title XIX” tag, a reasonable guess is that the vast majority of them are dependents; Medicaid legislation between 1984 and 1990 predominantly dealt with expanding coverage to children. Accordingly, we treat other Title XIX recipients as children (under age 21).

Medicare. Medicare predominantly serves the elderly but, since 1973, has been available to SSDI recipients. State-level data on Medicare expenditure per disabled enrollee were rarely published by the HCFA outside of a brief period from 1980 to 1983. Fortunately, the Census published these data in many of their own reports in the 1980s and early 1990s. The HCFA estimates were an input into Census' calculations of the poverty rate, though the Medicare estimates were not subsequently included in the microdata files of the CPS until much later (2004). With the exception of a single year (1988), Census' reports provide data on Medicare expenditure per disabled enrollee between 1982 and 1992. We linearly interpolate to impute 1988 estimates.

After 1992, we have to derive estimates of Medicare spending per *disabled* enrollee by using data on spending per enrollee among *all* Medicare participants. The source of data on state-level average spending among all enrollees is the State Health Expenditure Accounts (SHEA), developed and maintained by the Office of the Chief Actuary at CMS.[46] We consider two ways of inferring spending per disabled enrollee from the SHEA. For one, we calculate for each year following 1992 the ratio of spending per disabled enrollee in the U.S. as a whole to spending per enrollee among all Medicare participants in the U.S. We then multiply Medicare spending per enrollee in each state by this aggregate ratio. In a second set of estimates, we make state-by-state adjustments for disabled-enrollee spending, but the adjustment factors are based on the early period 1980 to 1992. Using the Census-published HCFA data, we first calculate the ratio of spending per disabled enrollee in each state to spending per enrollee among all participants, and take an average over this early period. Then, to derive an adjustment factor for each state in each year after 1992, we scale all states' ratios by the same factor so that total U.S. spending per disabled enrollee implied by our state-level panel equals the published estimate from the CMS in each year.

The construction of the SHEA and Census-published estimates differ in a few details. First, the Census does not excise the portion of the Medicare premia paid by enrollees, whereas the SHEA does so (and thus reflects only Medicare reimbursements for care).

[45] Title XIX refers to the section of the Social Security Act that concerns Medicaid.

[46] Levit et al. (1995) introduced the SHEA data.

Second, the SHEA estimates are adjusted for border-crossing, which means spending is assigned to states where the enrollees live rather than where care is provided. The Census-published estimates do not appear to do this.

Accordingly, we are inclined to believe that the aggregate spending levels in the SHEA are more reliable. Therefore, we splice together the (adjusted) SHEA and Census-published estimates by extending the former back in time using the growth of the latter.

Market values. In the main text, we noted that Medicaid and Medicare expenditure per enrollee is, under certain circumstances, akin to the premium in private insurance markets. Hence, Census refers to per-enrollee expenditure on these programs as the "market value" of public insurance.

However, this analogy with private markets can be taken only so far. In the case of private insurance, premia are revealing insofar as they signal at least a lower bound on the willingness to pay. It is less clear that expenditure per enrollee signals the same information. In other words, the market value can be a poor measure of recipients' valuation of Medicaid.

Elwood and Summers (1985) are forceful advocates for this view. They argue that Medicaid should be considered income—and hence valued as a means to raise consumption—only to the extent that it substitutes for what a household would otherwise *choose to pay* for health insurance. If the latter is zero, the "fungible value" of Medicaid is said to be zero—even for an enrolled household. This argument underlies the Census' view of Medicaid's fungible value, an estimate of which is included in the CPS data files as an alternative to the market value.

There are at least three challenges to the fungible-value approach, however. First, even if we grant that a poor household's marginal utility of, say, living space exceeds that of health insurance, it need not follow that Medicaid should be ignored entirely as a source of income. The latter view can be defended only in the extreme case of an elasticity of substitution between living space and insurance is zero (Weicher 1999). Second, the fungible-value approach is difficult to implement insofar as one must take a stand on what a household would have purchased in the absence of Medicaid. The Census' approach is to estimate "required" expenditure on basic necessities, such as food and living space. If an enrolled household's income exceeds this required spending, it infers that the household would purchase private insurance in the absence of Medicaid. Determining what is "required" spending, however, is fraught with difficulty (Weicher 1999). Finally, the fungible-value approach is ill-suited to parameterizing the value of *means-tested* insurance. Suppose a household chooses not to work, and thus earns very little income, *in order to* retain access to Medicaid. One would not want to infer that its value of Medicaid is zero.

Unemployment insurance. The TRIM program does not consistently correct for under-reporting of UI receipt. In its place, we have devised a simple imputation algorithm.

First, we identify the subsample of CPS respondents reporting both positive UI income and earnings. Dividing UI income by weeks of layoff yields an estimate of the weekly benefit amount; dividing earnings by weeks of work yields weekly earnings.

Now, letting $\ln b_{ist}^{UI}$ denote the log of weekly UI income of individual i in state s and year t and $\ln w_{ist}$ the log of weekly earnings, we specify a regression,

$$\ln b_{ist}^{UI} = \alpha_s + \psi_s \ln w_{ist} + \varepsilon_{ist}, \tag{8}$$

where α_s is a state fixed effect. Equivalently, by expressing UI income and earnings relative to their (CPS-implied) state specific means, we can write this as

$$\ln b_{ist}^{UI} - \overline{\ln b_{st}^{UI}} = \psi_s\left(\ln w_{ist} - \overline{\ln w_{st}}\right) + \xi_{ist}, \tag{9}$$

where $\overline{\ln b_{st}^{UI}}$ is the mean over i of log weekly UI income; $\overline{\ln w_{st}}$ is the mean over i of log weekly earnings; and $\xi_{jst} \equiv \varepsilon_{jst} - \overline{\varepsilon_{st}}$. In principle, this regression model can be estimated separately in each year, but annual CPS samples are quite small in several of the states. Therefore, we pool data across years.

Next, we impute *potential* weekly UI income to every worker reporting any weeks of unemployment. This involves two steps. First, recalling that our regression is specified in log deviations from means, we calculate an estimate of UI weekly income, $\tilde{b}_{ist}^{UI}$, equal to the sum of (i) $\exp\left[\hat{\psi}_s\left(\ln w_{ist} - \overline{\ln w_{st}}\right)\right]$ and (ii) the average weekly benefit according to the Employment and Training Administration's (ETA) Handbook 394 report. We use these administrative data for the mean benefit at this stage rather than the CPS measure ($\overline{\ln b_{st}^{UI}}$) due to concerns over reporting error in the latter. [47] In the second step, we take account of the statutory caps on UI weekly benefits in each state. Accordingly, we assign a potential weekly UI benefit $\hat{b}_{ist}^{UI}$ equal to the minimum of $\tilde{b}_{ist}^{UI}$ and the state's maximum benefit.[48] Thus, a respondent's potential annual UI income, $\hat{B}_{ist}^{UI}$, is given by the product of $\hat{b}_{ist}^{UI}$ and his reported number of weeks of unemployment.

Of course, a substantial share of potential UI recipients do not take up the benefit [see Blank and Card, 1990]. The last step in our imputation algorithm, then, is to select a sample of UI recipients from the set of potential participants. To this end, we first calculate the ratio of ETA's report of actual benefits paid by state s in year t to the sum of $\hat{B}_{ist}^{UI}$ over all potential recipients in the CPS in state s in year t. Denote this ratio by ϕ_{st}. A potential recipient i is then included in our sample of UI participants if $\phi_{st} >$

[47] This makes only a minor difference, since the principal problem is underreporting of receipt, rather than misreporting of b_{ist}^{UI}, which refers to weekly UI income *among those* reporting receipt. To see this, we compare (i) total benefits paid according to the ETA to (ii) a counterfactual estimate equal to the product of average annual UI income in year t among those reporting in the CPS and the total number of workers who began their UI spell in t. The latter, taken from the ETA, is a lower bound on the number of individuals who ever received UI in t. The counterfactual (ii) is typically within 10 percent of (i).

[48] Data on maximum benefits is available at https://ows.doleta.gov/unemploy/statelaws.asp#sigprouilaws.

$u_{ist} \sim U[0,1]$. By construction, then, total UI income in our sample of participants will replicate the ETA's report.

C. Results based on the SIPP (in progress)

The SIPP was designed in large part to measure eligibility for and participation in public assistance programs. It is a longitudinal survey that follows a panel of households for up to four years. The number of households in any panel is similar to the size of CPS samples. The Census has completed surveys of six panels since the SIPP was first fielded in 1984.

While SIPP tends to understate benefit receipt on average, the *change* in under-reporting is typically statistically insignificant (Meyer et al, 2015). Indeed, in the case of SNAP/Food Stamps, Meyer et al. find virtually no change in under-reporting. They estimate an increase in under-reporting of AFDC/TANF and SSDI, but the standard errors on the time trends are too large to rule out no change whatsoever. One exception to this theme is UI, which has exhibited a marginally statistically significant increase in underreporting.

For our purposes, an important drawback of the SIPP—and the reason we do not use it as our baseline source of data—is that it was first administered in 1984. Nonetheless, we can compare our evidence based on the CPS with SIPP estimates over this latter period. The relative reliability of the SIPP in terms of its measurement of benefit income makes it a natural cross-check on our CPS estimates.

Though the SIPP is a panel, we organize it into a series of repeated cross sections to mimic the structure of the CPS that we have used thus far. In each calendar year, we identify SIPP respondents who participated in each of that year's three interviews. Each interview asks about weeks worked and income received in a four-month stretch of the year. Thus, for these households, we can construct a record of annual weeks worked and income analogous to what we observe in the CPS. We then re-run the income gradient regressions specified in section 2.1 on the repeated cross sections constructed from the SIPP. The results are still quite preliminary but worth mention.

Unlike in the case of the CPS, the SIPP results appear more sensitive to the regression specification. If we use the log of household income, the income gradient declines considerably more than in the CPS. In the SIPP, the gradient falls by 46 log points between 1987 and 2013, whereas it declines by 14 in the CPS. [49] A closer look reveals that the discrepancy is largely due to a rather remarkable 20 log point decline in the SIPP in the last half of the 1990s, which is, notably, a period when nonemployment actually fell. If we use the *level* of household income, the change in the gradient across the SIPP and

[49] The 2008 panel was completed in 2013. The panel initiated in 2014 is still in progress.

CPS is more similar: the gradient falls 23 log points in the SIPP and 13 in the CPS.[50] The difference across log and level specifications in the SIPP appears due to a relatively large increase in the variance of wage income among the full-year employed. By Jensen's inequality, this amplifies the difference in the level of income across weeks relative to the difference in the log, which buoys the gradient when the latter is estimated in levels.

[50] The gradient in each year in this case is calculated as the ratio of the regression coefficient to average household income.

Appendix references

Cooper, Daniel H., Byron F. Lutz, and Michael G. Palumbo. 2011. "Quantifying the Role of Federal and State Taxes in Mitigating Income Inequality." Federal Reserve Bank of Boston Research Department Public Policy Discussion Papers, No. 11-7.

Elwood, David and Lawrence Summers. 1985. "Measuring Income: What Kind Should Be In?" Proceedings of the Conference on the Measurement of Noncash Benefits, U.S. Bureau of the Census.

Levit, Katharine R., Helen C. Lazenby, Cathy A. Cowan, Darleen K. Won, Jean M. Stiller, Lekha Sivarajan, and Madie W. Stewart. 1995. "State Health Expenditure Accounts: Building Blocks for State Health Spending Analysis." *Health Care Financing Review* 17(1): 201-254.

Weicher, John C. 1999. "Some Income-Measurement Issues and Their Policy Implications." *American Economic Review Papers and Proceedings* 89(2): 29-33.

Winterbottom, Collin, David Liska, and Karen Obermaler. 1995. *State-Level Databook on Health Care Access and Financing*. Washington, D.C.: The Urban Institute.